

Roman Drahtmüller, Viviane Glanz, Roland Haidl, Jana Jäger, Jordi Jaen Pallares, Karine Nguyen, Edith Parzefall, Peter Reinhart, Marc Rührschneck, Thomas Schraitle, Martin Sommer

SuSE Linux Connectivity Server

Installation – Configuration – Administration



SuSE GmbH
Schanzäckerstr. 10
D-90443 Nürnberg
Phone: (09 11) 7 40 53 520 (Enterprise Sales)
(04 21) 5 26 23 00 (Support; for hours, see the documentation)
e-mail: suse@suse.de
WWW: <http://www.suse.de>

Roman Drahtmüller, Viviane Glanz, Roland Haidl, Jana Jäger, Jordi Jaen Pallares, Karine Nguyen, Edith Parzefall, Peter Reinhart, Marc Rühschneck, Thomas Schraitle, Martin Sommer

SuSE Linux Connectivity Server

1st Edition 2001

(c) SuSE GmbH

Copyright

This product is the intellectual property of SuSE GmbH.

It may only be copied in its entirety or in excerpts only if each copy is marked with this copyright label.

Layout: L^AT_EX 2_ε

Geeko icons by Rolf Vogt

English Translation: Rebecca Ellis

Linux is a trademark of *Linus Torvalds*. XFree86TM is a trademark of *The XFree86 Project, Inc.* *Windows*, *Windows 95*, *Windows 98*, *Windows ME*, *Windows NT* and *Windows 2000* are registered trademarks of the *Microsoft Corporation*. *UNIX* is a registered trademark of *X/Open Company Limited*. Other (registered) trademarks included herein are: *T-Online* by *Deutsche Telekom*, *SuSE* and *YaST* by *SuSE GmbH*. All product names are used without guarantee of free usability and may not include trademarks. The corporation SuSE GmbH essentially conforms to the written format of the vendor. Other products named here can be trademarks of their respective vendors.

Contents

Foreword	vii
1 Why Linux?	1
1.1 The Alternative Called Linux	1
1.2 The Technical Side of Things	1
1.3 The Philosophy Behind It	4
1.4 Conclusion	6
2 Support and Services	7
2.1 No Product Support or Maintenance without Registration	7
2.2 Support Services for SuSE Linux Connectivity Server	7
2.3 Product Support	7
2.4 Maintenance für den SuSE Linux Connectivity Server	8
2.5 The fastest way to help!	9
2.6 How Do I Reach the Support Team?	10
3 Installation with YaST2	11
3.1 Starting Your System from CD-ROM	11
3.2 The Opening Screen	11
3.2.1 Other Installation Options	12
3.3 YaST2 Takes Over	13
3.4 Selecting a Language	14
3.5 Mouse Pointer	14
3.6 Keyboard and Time Zone	15
3.7 Selecting the Hard Disk	15
3.7.1 Selecting Partitions	16
3.7.2 Note for Advanced Partitioning	17
3.8 Logical Volume Manager (LVM)	19
3.8.1 Configuring LVM with YaST2	19
3.8.2 LVM – Partitioning	19
3.8.3 LVM – Setting Up Physical Volumes	20
3.8.4 Logical Volumes	22

3.9	Configuring the Crypto File System	23
3.10	Boot Manager for System Start-Up	24
3.11	Root Password	25
3.12	Creating an Administrator Account	26
3.13	Let's Go!	27
3.14	Preparing the Hard Disk	27
3.15	Installation of Software Packages	27
3.16	Monitor Settings	28
3.17	Network Card	30
3.18	Finishing the Installation	32
3.19	Graphical Login	32
4	SLCS Server Configuration with YaST2	35
4.1	Hardware	36
4.1.1	Printer Configuration	37
4.1.2	Graphical Interface (X11)	39
4.1.3	Keyboard	40
4.2	Internet Access	40
4.2.1	Basic Internet Connection	40
4.2.2	Instructions for all Types of Internet Access	42
4.2.3	Internet Connection and Local Network	42
4.2.4	ISDN	43
4.2.5	Modem	44
4.3	Network Card	45
4.4	Server Services	48
4.4.1	Basic Samba Configuration	48
4.4.2	NFS Server Configuration	48
4.4.3	NIS – Network Information Service	50
4.4.4	E-mail (Sendmail)	52
4.5	System	53
4.5.1	Managing Users and Groups	53
4.5.2	System Security	55
4.5.3	Install and Remove Software	57
4.5.4	Change Installation Source	58
4.5.5	Online Update	59
4.5.6	System Update	61
4.5.7	Boot Mode	61
4.5.8	Creating a Boot, Rescue or Module Disk	62
4.6	Miscellaneous	63

4.6.1	Hardware Information	63
4.6.2	Start Protocol	64
4.6.3	System Protocol	65
4.6.4	Loading the Vendor's Driver CD	65
4.6.5	Creating Backups	66
4.7	Important Variables in the rc.config Editor	67
5	Workstation Configuration	69
5.1	Access to the File Server	69
5.1.1	Windows 95/98/ME	69
5.1.2	Windows 2000	72
5.1.3	With MacOS	75
5.2	Internet Access	76
5.3	Configuration of SuSE Linux Clients	77
6	Network Services – Behind the Scenes	81
6.1	Basic Functions	81
6.1.1	Domain Name Service	81
6.1.2	DHCP	83
6.1.3	NIS	84
6.2	File and Print Service	85
6.2.1	NFS – Shared File Systems	85
6.2.2	Samba	87
6.3	Security	88
6.3.1	Firewall	88
6.4	Proxy Server: Squid	90
6.5	Intranet Server	93
6.5.1	Apache	93
7	ALICE	97
7.1	What is ALICE?	97
7.2	How ALICE is Installed	97
7.3	Creating a Simple Configuration	98
7.4	Further Information	100
8	Security and Confidentiality	101
8.1	Basic Considerations	101
8.2	Local Security and Network Security	101
8.2.1	Local Security	103
8.2.2	Network Security	107
8.3	Some General Security Tips and Tricks	110
8.4	Using the Central Security Reporting Address	112

9	Troubleshooting	113
9.1	Creating a Boot Disk	113
9.1.1	Creating a Boot Disk In DOS	113
9.1.2	Creating a Boot Disk with UNIX	114
9.2	LILO Problems	115
9.2.1	Diagnosis of Errors: LILO Start Messages	116
9.2.2	The 1024-Cylinder Limit	117
9.3	The SuSE Rescue System	119
9.3.1	Working with the Rescue System	120

Foreword

The SuSE Linux Connectivity Server is the ideal tool for small commercial networks designed for small business without an on-site system administrator. This server offers you an all-in-one solution for your clients' Internet and Intranet: file server, print server, Internet gateway with proxy, and firewall.

This is accomplished by the SuSE Linux Connectivity Server, following a relatively simple installation procedure, and does not necessitate complication system maintenance.

If you have special requirements, extending beyond the preconfigured settings, you can obtain support services from SuSE Professional Support.

The SuSE Linux Connectivity Server, which is quite robust as compared to other products, features long release cycles. It combines the proven stability of a Linux system with constant up-to-dateness — Online Updates can be automatically installed at any time. Moreover, this will spare you an unnecessary burden on your finances.

SuSE designed the SuSE Linux Connectivity Server in order to ensure utmost productivity and profitability in terms of your business network, so that you are free to concentrate on the essentials.

Have a lot of fun !

Your SuSE Team

1 Why Linux?

1.1 The Alternative Called Linux

Linux stands out from other established operating systems in terms of its unique product philosophy. The principle of open source and free distribution, the success of the open developer's model, and the integrated technical features constitute Linux's strengths.

1.2 The Technical Side of Things

This section will focus on a few technical aspects, which distinguish Linux from other systems.

The Kernel

The innermost core of the Linux system – the kernel – harbors some of the most key secrets to this operating system's success.

Comprehensive hardware support The kernel is laid out in such a way that it can support practically the entire spectrum of available hardware, from the smallest hand-held to the mainframe. Thanks to open source, it is relatively easy to port Linux to new hardware.

Network capability Since its “childhood days”, Linux's main focus has been on network capability. Through its TCP/IP protocol support and its support of other open Internet standards, networked communication was already an option even in Linux's early developmental stages. Once again, it was its network capability that gave Linux the leverage it needed for further development – Linux was and is continuing to be developed over the Internet.

Security The development of system security gates also accompanied network capability. After all, the personal computer also has to be guarded from attacks originating from the Internet. Linux already provided support for IP filtering, even in its early developmental stages. In any case, the original code has, time and time again, undergone rigorous restructuring to ensure that it remains compatible to the constantly changing face of the Internet.

Performance Linux supports some essential “tricks” for dealing with system resources designed to significantly improve its performance in comparison to other systems. For example, Linux automatically generates dynamic hard

disk data caches during operation. In this vein, it works in “read ahead” mode (provisionally reading sectors in advance) and “delayed write” mode (reserves write access for execution these permissions in one go). The “delayed write” procedure is also the reason why you should not just switch off a Linux machine. Both of these aspects are responsible for the main memory only seeming to fill up with time and are the reason why Linux is so fast as well.

In particular, as of Version 6.3, SuSE Linux features full support for LVM (Logical Volume Manager). As of Version 7.2, the LVM can even be configured in an installed system with the help of its own YaST2 modules. This option is especially useful for those who working with memory management of a large scale, such as for databases.

Multituser and Multitasking

Just like its UNIX ancestors, Linux is a real multiuser system. That means that a system can be used simultaneously by multiple users. They can be working directly from affiliated terminals — or, and which is usually the case — access it over the network. In contrast to Windows (NT) environments, where the server must communicate with the client applications on the client itself, the user logged in to Linux has access to a complete user environment with all its services.

Even for stand-alone workstations, which do not have access to a network and where only one user is working, multiuser capability has real advantages. Even the normal desktop station features additional virtual consoles, along with a graphical interface. If, for some reason, the graphical user interface is not responding, you still have the option of switching to a virtual console and logging in to a shell from there, and then, restarting the graphical user interface. This way, rebooting is practically unnecessary since usually, in Linux, only one program at a time might crash, but not the entire operating system all at once. This aspect leads to the next clincher for multitasking capability.

In Linux, several processes can be running at simultaneously. The operating system has direct control over the processes and decides when a process is suspended. This design improves system performance, even if you do not have two actual parallel-functioning processors.

Security

Linux’ system design and that of its UNIX ancestors automatically entails some important security advantages in comparison to to other systems.

As a multiuser system, Linux is designed to support multiple users working simultaneously on one machine, but where no normal user has full control over the system. This way, no user can cause harm to any other user environments by corrupting the entire system. Which user and which group a given directory, process, or file belongs to is strictly regulated. Only ‘root’ has access to the system resources.

Due to this reason, it is significantly more difficult for computer viruses to cause damage Linux systems. A virus inside a mail attachment can never access the system resources and thereby afflict the entire system:

- The user lacks the file permissions to run a program in an attachment such that it could affect any changes to the entire system.
- A virus can only cause damage to a system if the virus obtains root access by (carelessly) being run by superuser `'root'` or by taking advantage of security gaps.
- There are dozens of Linux mail programs — it is more difficult to cause damage across the board in a mixed environment than in a “monoculture”.

In addition, security gaps in programs whose code could potentially be read by anybody on this planet, are detected and fixed more quickly than in software whose sources are not freely available. There, potential security gaps may linger without ever being detected.

Free software, on the other hand, is transparent and anybody who finds a security gap can draw attention to it. To this end, there are even special mailing lists, web sites and forums such as *SecurityFocus* (<http://www.securityfocus.com>). Usually, the developers of the affected software — often in collaboration with those who have discovered the security gap — can quickly solve the problem. Frequently, what are known as “code audits” are instituted. Project developers will then critically examine the code of another developer group.

However, freely accessible program source codes can also give computer criminals the opportunity to track down security gaps and take advantage of them.

But, all in all, the positives of open source outweigh the negatives, since the respective codes are getting better all the time under the close scrutiny of numerous co-programmers, as well as in terms of continuing development.

A Wide Variety of Applications

Some of the greatest achievements of open software is the sheer volume of applications available to an unlimited number of users at no further cost.

There is not just one standard solution to meet whichever needs may arise, but a wide array — you, as user, have the freedom to pick the alternatives which work best for you. You can, of course, also participate in optimizing existing solutions by contributing your ideas and programming skills. This degree of freedom and flexibility was previously unknown to the world of software.

Switching to Linux in larger companies can still be problematic for larger companies, due to the continuing lack of *open source* solutions for certain specialized types of applications or, due to existing programs not being able to function in Linux. Met with increasing acceptance in a variety of industries as well as by the public, it is only a question of time until these alternatives do exist and commercial solution providers port their software to Linux. The most recent example of this is Oracle’s Linux offensive.

Support

The usage of free software is often discouraged since no professional support exists in this sector. These reservations can be addressed in two different ways:

An e-mail to the respective supporter mailing list can provide short-term free assistance. Anyone, regardless of whether they are an individual user or a company client, can go through these channels relatively unhindered by red tape. However, with the increasing success of free software, the ratio of users to developers is changing in such a way that individual developers are finding it increasingly difficult to respond to requests.

This is why this more traditional method is being supplemented by professional services specializing in the support of free software products. These are not only the Linux distributors (SuSE, Red Hat, Caldera, etc.), independent companies and small, specialized companies, but also, more and more frequently, large companies such as IBM, which is in the process of evaluating free software's potential. Here, companies in need of support can purchase it and can even reach the corresponding developer if they have dire support needs.

1.3 The Philosophy Behind It

Open Source

At the very beginning of UNIX history, the tradition of distributing software simply by passing along the source code was born. The recipient took the code, adjusted it to his hardware requirements, and consequently obtained an executable operating system tailored to his own hardware. During the continuing course of the commercialization of the UNIX market, this tradition was slowly forgotten. One of the few exceptions to this phenomenon is Richard Stallman's GNU project.

Linux represents a sort of reawakening of this "tradition". Open source allows program code to be easily ported to new platforms. Anyone can take the Linux source code and do what he deems sensible with it. He does not have to just put up with any software errors, but can fix them himself and thus contribute to improving the code's foundations as well as profit from the contribution of other like-minded folks.

Open and free source code has an immense significance for a successful development process. The following points illustrate just why the *open source* principle is so important:

- Significantly more developers and testers exist for open source as do for a closed project.
- Although several programmers at a time are working on one product, regular communication, strict control mechanisms, and a small coordinator team, all ensure that "too many cooks are not spoiling the recipe".
- The open availability of the source code and the large number of potential developers lessens the dependency on just one contact person.

- Tracking and fixing bugs is faster and more efficient.
- There is better feedback between users and developers.
- Decisions on design are openly discussed. Errors can be detected in a timely manner and eradicated.

The consistent use of the Internet contributes to the enormous progress in development and to the world-wide propagation of the system. Developmental decisions are discussed on mailing lists and over special Internet forums. The various software versions are managed in CVS and are consistently maintained. The newest versions, including bug fixes and updates, are distributed world-wide over FTP servers.

Free Software

The definition of “free” software extends beyond just the open availability of the source code at no cost. Strictly speaking, the concept of *free software* entails more — and is defined in the *GNU General Public License (GPL)*.

“Open source” also implies the freedom to do what you want with the software, and also to modify and adjust it according to your own needs, not to mention, to pass the source code along.

To protect this freedom, the *GPL* contains restrictive clauses guaranteeing the open availability of *GPL* code. Thus, any code derived from *GPL* code must ultimately fall under the definition of *GPL*. The program source must be publicly accessible and remain available free of purchase. Those who pass along this source code may not receive compensation for doing so, apart from material costs.

Motivation

Sooner or later, anyone who is grappling with the phenomenon of *open source* will ask herself what could motivate these countless developers to push free software. The traditional idea of wage labor does not apply here, of course.

Some motivating factors for working on *open source* projects are outlined in the following:

- Due to the sheer number of open software programmers, each one benefits from work done by the others. She receives more in return than she could program herself.
- The collaboration on a free software project offers — as in the academic world as well — the opportunity to gain good standing and respect among like-minded people.
- In the *open source* sector, software technology belongs to everyone. *Anyone* can join in and do what they want with the code. Previously, this was only the privilege of a few companies.

- The programmer can develop free software with completely free developer's tools (compilers, developer's environment). This extends the developer community to many programmers who might otherwise not have had the means to participate in the development process.

Many companies, whether they are Linux distributors or consulting firms, smaller or larger hard- and software manufacturers, take advantage of this motivating potential and either directly or indirectly finance development on open software projects.

1.4 Conclusion

The strengths of an *open source* solution are obvious. The software critical to your enterprise will be running and stable — since the software has already been scrutinized by countless testers by way of the open source developer model.

Security-related issues are expediently remedied and the related updates are made available over the Internet.

Open source software supports a wide variety of hardware. This way, you do not have to depend on the manufacturer and are even able to get relatively low-performance, older hardware working for you, providing a longer useful life for your hardware and the relative independence you will enjoy overall from innovation cycles will turn out to be an excellent advantage.

Open source is based on open standards. If you are, for instance, dependent on long-term data archiving, you will benefit from the use of non-proprietary, manufacturer-independent file formats. By supporting open and standardized communication protocols, your Linux system will easily master data exchange with other systems — regardless of operating system and manufacturer.

If any one of these points mentioned above are areas of concern for you, you will definitely be right on the mark by choosing an *open source* solution.

2 Support and Services

2.1 No Product Support or Maintenance without Registration

In order to be able to guarantee the best product support possible, only requests from registered users will be responded to. On the back of the CD cover, you will find two stickers, each labeled with a product registration code. This code is unique and serves to verify whether you are authorized to receive these services.

When you send us your data in an online form at

<https://support.suse.de/en/register/>

you can become a registered user with access which includes product support and maintenance.

You can also send us your registration with the enclosed registration card by mail. To do this, affix a registration code sticker in the designated field on the registration card. We recommend that you leave the second label on the product so that you can always have the registration code at hand if you need to consult our support services.

2.2 Support Services for SuSE Linux Connectivity Server

The product support already included in the purchase price of the SuSE Linux Connectivity Servers is good for a period of 30 days after the registration date and covers the services listed below.

This support is not intended as training material nor as an introduction to SuSE Linux itself, but as a guide for the basic installation of the system. Support can therefore be requested only in respect to configuration problems, not pertaining to conceptual questions.

2.3 Product Support

Product support covers the basic installation of the SuSE Linux Connectivity Servers on a machine, as well as the configuration of the basic hardware and the following peripherals, using the configuration tool YaST2:

- Local Printer (over lpd)
- Network Card (Ethernet)
- Modem
- ISDN
- DSL

Product support also provides you with support in the configuration of the network services listed below, with YaST2:

- Windows and Apple network drives (samba and netatalk)
- Proxy server for Minimizing Internet Traffic (squid)
- Mail relay (sendmail)
- Central user management for UNIX clients over NIS
- Automatic Internet dialup if necessary
- Support for configuring the following services on your clients:

SuSE Linux: DHCP, DNS, NFS, NIS, NTP, Proxy, Samba

Windows: DHCP, DNS, Proxy, Samba

Mac: AppleTalk, DHCP, DNS, Proxy

2.4 Maintenance für den SuSE Linux Connectivity Server

The maintenance of the SuSE Linux Connectivity Server III is an active maintenance contract, preventative support, customized according to your specific IT requirements. You will receive the following services, which is up-to-date and which guarantees utmost user-friendliness:

SuSE Linux Connectivity Server Maintenance is an active maintenance contract, preventative support, which will meet your highest standards. You will receive the following services, guaranteeing that your technology is always current and easy to use.

- Fixes and patches for resolving critical errors (security, data loss) of the SuSE Linux Connectivity Servers.
- Each patch includes thorough documentation.
- You will be contacted by SuSE Enterprise Support Services by mail.
- The patches themselves will be made available on a secure web server for downloading.

- You will obtain support for installing patches from SuSE Enterprise Support Services.

You also have the option of optimizing your maintenance service with an update service. In addition, you will receive all patches and fixes for your SuSE Linux product quarterly, shipped on CD.

Your registration automatically entitles you to SuSE Linux Connectivity Server Maintenance for a period of 12 months. Thus, you will be guaranteed a stable and tested system.

2.5 The fastest way to help!

Register your product online at our web site: <http://support.suse.de/en/> and send your request by e-mail to slcs-support@suse.de.

Please give your customer information in the e-mail text before describing the problem. Note the usage capitalization and lower-case spelling in your customer information. This way, your e-mail can be automatically processed (see Example 2.5). Do not use any unnecessary attachments and, if you need to insert configuration files, enter them directly in ASCII format in the request letter.

Example:

```
FIRSTNAME: John
LASTNAME: Doe
COMPANY: Doe & Co. Inc.
STREET: Easy Street 7
CITY: Nowhereville
ZIP: 12345
COUNTRY: USA
REGCODE: <Product registration code>
EMAIL: doe@doe-inc.com

My Problem: problem description ...

My Hardware: hardware description ...

<doe@doe-inc.com>
```

File 2.5.1: Support Request by E-Mail

2.6 How Do I Reach the Support Team?

You can reach the Support Team via the following contact information and during these hours of operation:

- E-mail: slcs-support@suse.de
Processing: weekdays
- WWW (e-mail): <http://support.suse.de/en>
Processing: weekdays
- Phone (calls answered by Enterprise Support Services):
Phone: +49 (0) 421 526 23 40
Open: Monday through Friday 9:00 a.m. – 6:00 p.m. (except for legal holidays)
- Fax: +49 (0) 911 74 05 34 77
Processing: weekdays

More information on our extended support services can be obtained at:

<http://support.suse.de/en>

3 Installation with YaST2

On the following pages, you will find instructions for installing SuSE Linux Connectivity Servers with YaST2.

This first chapter focuses on the following topics: starting the installation system, graphical installation with YaST2, hard disk partitioning, configuration of the boot mode along with the graphical interface, and a basic network configuration.



Note

The SuSE Linux Connectivity Server is immediately ready to be implemented following installation. All key server services are activated using default values so that no further configuration steps must be carried out under normal circumstances. If other servers are already up and running on the local network, conflicts could result. In this case, we recommend that you configure SuSE Linux Connectivity Server independent of your local network.

3.1 Starting Your System from CD-ROM

To start the installation, switch on your computer and insert the first SLCS-CD into the drive.

In order to be able to proceed with the installation, your system must be bootable from CD. If this is not the case, you might have to change your BIOS settings or, if SCSI systems are being implemented, the boot sequence of your SCSI controller.

Then consult your manufacturer documentation.

If your machine does not boot from the CD-ROM, you will have to change the settings in the computer's BIOS depending on what kind of CD-ROM drive is in the machine. You can find more information on this in Chapter 9.1 on page 113.

3.2 The Opening Screen

A screen as in Figure 3.1 on the following page shows you that the system is ready to be booted for the installation. Be sure to select 'Installation' here (the default selection). Either wait a few seconds or just press **Enter** to load the kernel.

A few seconds later, a minimal Linux system is loaded which takes over the rest of the installation procedure. A number of messages and copyright notices will then appear on the screen. At the end of the loading process, the YaST2 program will start, and a few seconds later, the graphical interface of YaST2, the SuSE Linux installation program, will be displayed.




Figure 3.1: The Opening Screen in SuSE Linux

3.2.1 Other Installation Options

If you press any key before the wait time is up, automatic startup will be disabled, whereby you can take your time in selecting other options. These are especially useful given the default settings, if problems exist in the graphical display. As the actual launching of the installation to the hard disk is initially preceded by some dialogs and specific queries, you can always cancel in case there are problems and then choose different options following a reboot.

A Different Graphics Mode for YaST2

Choose the standard VGA (640x480) graphics mode compatible to any graphics card using the function keys. In the worst case scenario, you can also select pure text mode.

In the text-mode YaST2 screen, skip from one menu item to the next using the **(Tab)** key and inside a menu, using the **(↑)** and **(↓)** keys. **(←)** takes you to the next screen.

Other Ways to Install

Select other systems with **(↑)** and **(↓)**.

- If you choose 'Manual Installation', you have more options available to you, in particular, in selecting device drivers to be installed. However,

drivers will not be automatically loaded. This is normally only relevant for experts.

- A ‘Rescue System’ helps you safely start your computer if there are problems on your system. More information on the rescue system can be obtained in Chapter 9.3 on page 119.
- ‘Memory Test’ starts a very extensive memory test which takes quite some time to run through. It will, however, more accurately pin-point the memory error than the BIOS memory test when booting.

Now, by pressing **Enter** the selected system will be started.

3.3 YaST2 Takes Over

Now the actual installation of SuSE Linux starts with the YaST2 installation program. Figure 3.2 shows you what the screen will look like. During this phase, the hardware available on your system is checked and prepared for the installation. A bar in the middle of the screen shows the progress of the installation.

All YaST2 screens have a common format. On the left, help texts are shown, providing information on the current help topic. All entry fields, lists and buttons on the YaST2 screens can also be accessed by your mouse. If your cursor doesn’t move, your mouse has not been automatically recognized by Linux. You will then need to use your keyboard, as explained in the above section.

Following the language selection screen, you will be able to manually configure your mouse.



Figure 3.2: The hardware analysis

3.4 Selecting a Language

SuSE Linux and YaST2 are adapted to use the language you have selected. English is the default setting for the English distribution of SuSE Linux. These settings can be changed individually.

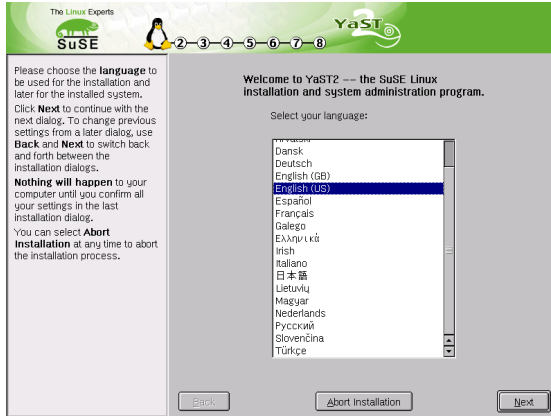


Figure 3.3: Selecting the language

If your mouse cursor still doesn't work, press the **(Tab)** key repeatedly until the 'Next' button appears, then press the **(←)** key.

3.5 Mouse Pointer

If YaST2 didn't recognize your mouse type automatically, an entry screen will appear as shown in Figure 3.4.

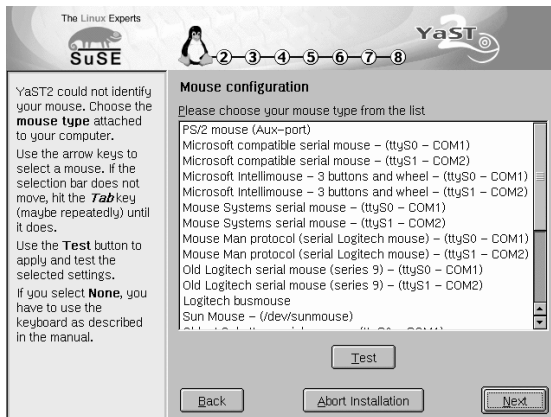


Figure 3.4: Selecting the mouse type

To select your mouse type, use the **↑** and **↓** keys. If you have documentation for your mouse, this should provide you with a description of the mouse type. Select the mouse type from the list. The first three items in the list are the most common mouse type. Try these first, if you don't know the type of your mouse. Confirm your selection either by pressing the **(Alt) + (T)** keys or pressing **(Tab)** and then confirming this with **(↵)**.

Now test if the mouse pointer on the screen follows your movements. If the cursor does not move, select a different mouse type and try again.

3.6 Keyboard and Time Zone

The next step (Figure 3.5) the keyboard layout and the time zone are selected. In the field 'Hardware clock set to', you can choose between local time and GMT. Your selection depends on the clock settings in the BIOS of your computer. If this is set to GMT, SuSE Linux will automatically apply the time change for Daylight Savings, Standard Time, and vice versa.

Now select the desired keyboard layout. Usually, this corresponds to the language you chose. Select the correct time zone in the other column. Test your keyboard with special characters such as **'|'** and **'@'** to see if they appear correctly on your keyboard. If it does not work, you chose the wrong layout. The installation will be continued with 'Next'.

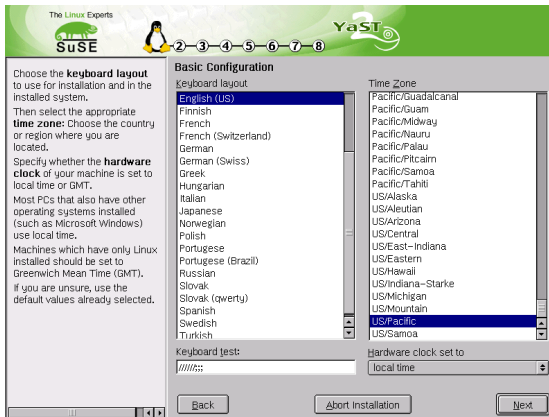


Figure 3.5: Selecting the keyboard layout and time zone.

3.7 Selecting the Hard Disk

Next, select the hard disk where the SuSE Linux System is to be installed. All the hard disks found on your system will be listed (see Figure 3.6 on the next page). Select the hard disk you want to install SuSE Linux Connectivity Server on.

Normally, the SuSE Linux Connectivity Server will be the only system on your machine and installation will only take place on a single hard disk. In this case, simply choose the hard disk then ‘Entire hard disk’. YaST2 will subsequently carry out the appropriate partitioning, whereby all data existing on the hard disk may be deleted, in order to make the entire disk space available for the SuSE Linux Connectivity Server.

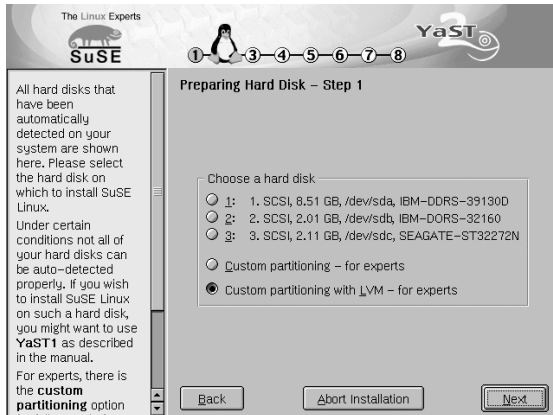


Figure 3.6: Selecting the hard disk where SuSE Linux is to be installed

If the system requires customized partitioning, however, or you want to use a logical volume manager, partition your own hard disk by selecting the option ‘Advanced settings, manual partitioning’, or ‘Customized partitioning with LVM - for experts’. Further information on configuring LVM can be found in Section 3.8 on page 19.

Note

Changes will not be applied to your hard disk until you have configured all installation settings and confirmed them in the designated dialog window with ‘Yes’. You can always return to the previous configuration screen to reset the changes you made while installing with YaST2, by clicking ‘Back’.

Default partitioning incorporates three primary partitions: a **boot** partition for the Linux kernel (approx. 20 MB) in the boot cylinder of the hard disk, a **swap** partition, fitted to the size of your RAM, and a **/** (or **root**) partition, for all system and user files, which take up the remaining hard disk memory.

3.7.1 Selecting Partitions

Once the hard disk, where SuSE Linux is to be installed, has been selected, YaST2 will list all the partitions located on the selected hard disk (Diagram 3.7 on the next page). Decide whether to ‘Use entire hard disk’ for SuSE Linux Connectivity Server and which partitions to delete in order to make more

room for the SuSE Linux System . Consult the YaST2 help guide to find out more about partition selection.

Caution

All data on the selected partition for installation will be erased. You will likewise lose all hard disk data if you select the menu item 'Use entire hard disk'!

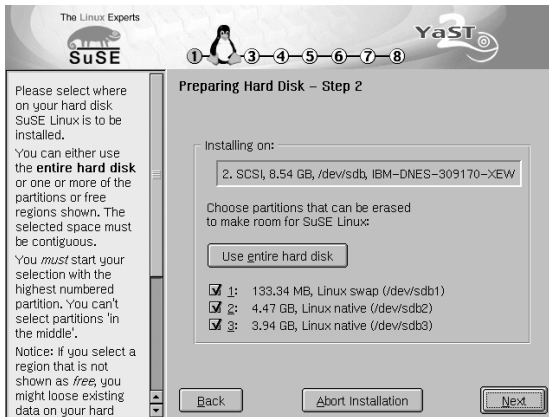



Figure 3.7: Selecting the partitions where SuSE Linux is to be installed

During the installation procedure, YaST2 will verify whether there is enough space on the hard disk designated for the SuSE Linux installation. If there is not sufficient memory, you will be prompted to make another selection. The installation of SuSE Linux Connectivity Server requires about 800 MB hard disk space.

3.7.2 Note for Advanced Partitioning

Only select this option if you are familiar with terms such as partitioning, mount-points or file systems. The default partitioning has already been configured for your system profile in YaST2. But still proceed with caution when partitioning your system.

You can 'Add', 'Edit', and 'Delete' partitions in this screen. See Figure 3.8 on the following page.

A suggestion for partitioning your hard disk might be:

/	2 GBytes
swap	double the RAM size, max. 1 GByte
/home	user directories; has its own directory under /home. Per user about 2 GByte.
/shared	shared directory, is accessible network-wide.

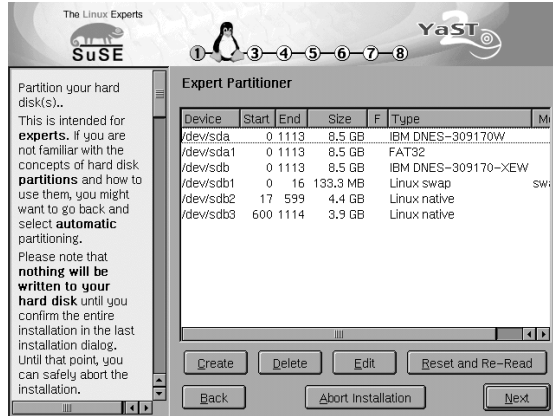


Figure 3.8: Selecting Partitions

If you want to expand the `/home` directory later, you can do this dynamically with LVM (see Section 3.8 on the next page).

The parameters for each partition on your system must be defined by hand:

- Define the size of each partition. Enter it directly in MBytes, or in hard disk cylinders.
- Decide on a format. You can choose between `ext2` or `reiserFS`. Choose `reiserFS` if you want to use the advantages of this journaling file system.
- Define a mountpoint for each partition. The mountpoint is the directory in your file system to which the partition is mounted. This option will most likely be useful if you want to store the `/home` or `/opt` directories, for example, on separate partitions.

Note

SuSE Linux Connectivity Server uses a shared directory for exporting common data for Samba and netatalk to all network clients on the system. If you want to use a separate partition for this shared directory, specify `/shared` as mountpoint.



3.8 Logical Volume Manager (LVM)

The Logical Volume Manager (LVM) enables flexible distribution of your hard disk space on several file systems. Since partitions can only be changed on a running system with relative difficulty, LVM was developed: this makes a virtual “pool” (volume group, or VG for short) of memory space available which can generate logical volumes (LV) as needed. The operating system will then access the LVs instead of the physical partitions.

Characteristics:

- Several hard disks/partitions can be merged into one large logical partition.
- If an LV (such as `/usr`) gets filled up, you can enlarge it given the appropriate configuration.
- You can even extend hard disks or LVs on a running system using LVM, provided that the “hot-swappable” hardware is suitable for such procedures, of course.

Using LVM is already quite beneficial for home PCs or small servers placed under high demand. If you have a growing data stock such as databases, MP3 archives, or user directories etc., the Logical Volume Manager might be just the right thing for you. With this, you could have file systems, for instance, which are larger than a physical hard disk. Another advantage of the LVM is that you can create up to 256 LVs. But please be aware that working with the LVM is quite different than working with conventional partitions.

Further information on configuring the “Logical Volume Manager” (LVM) can be found in the official LVM Howto:

<http://www.sistina.com/lvm/Pages/howto.html>

or at

`/usr/share/doc/howto/en/html/LVM-HOWTO.html`.

3.8.1 Configuring LVM with YaST2

You can activate the YaST2 LVM configuration by selecting ‘Custom partitioning with LVM’ while you are in the initial phase of preparing the hard disk, see Figure 3.6 on page 16.

3.8.2 LVM – Partitioning

First, you will reach a dialog where you can change the partitioning of your hard disk (see Figure 3.9 on the next page). If needed, add partitions and volumes. After clicking on ‘Add’, select the LVM type in the screen which follows, by clicking on ‘Do not format’ then specifying `0x8e Linux LVM` as ‘File System ID’. The LVM “partitions” do not need to be created yet. Therefore, you can ignore the warning which appears after clicking on ‘Next’. Also, keep in mind that no mountpoint has to be given yet. This is done at a later point.

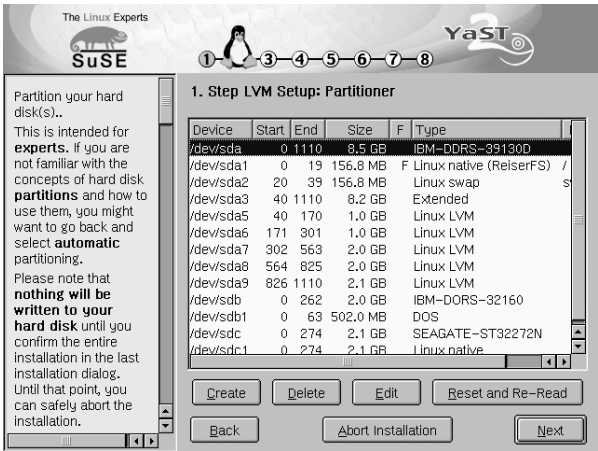


Figure 3.9: YaST2: LVM Partitioner

Note

In YaST2, at least the **Root** (or **/-**) file system must be located on a normal partition, such as on an **ext2** or a **reiserFS** partition.

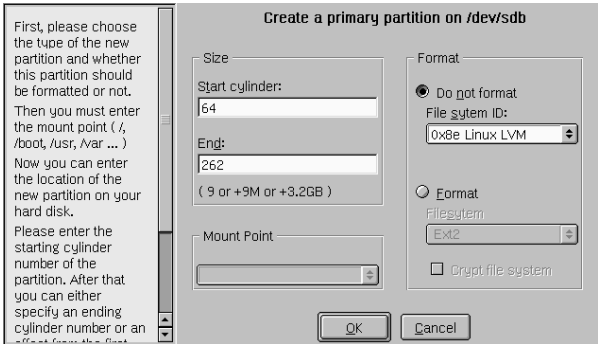


Figure 3.10: YaST2: Creating a LVM Partition

3.8.3 LVM – Setting Up Physical Volumes

This dialog manages the LVM volume groups (often abbreviated to “VG”). If there is no volume group yet on your system, you will be prompted by a pop-up window to create one. “System” is the name suggested for the volume group where your SuSE Linux system files are located. What is known as the physical extent size (often abbreviated to PE size) defines the maximum size of a

physical and logical volume in this volume group. This value is usually set to 4 megabytes. This allows for a maximum size of 256 gigabytes for a physical and logical volume. You should therefore only increase the physical extent size (e. g. to 8, 16 or 32 megabytes) if you need logical volumes larger than 256 gigabytes.




Figure 3.11: YaST2: Creating a Volume Group

In the following dialog, all partitions are listed that either have "Linux LVM" or the "Linux native" types. All swap and DOS partitions will therefore not be shown. If a partition is already assigned to a volume group, the name of the volume group will be listed. Unassigned partitions bear the label "--".

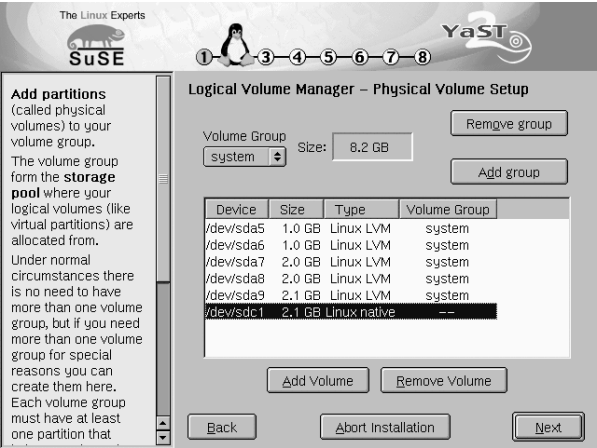


Figure 3.12: YaST2: Overview of the Partitions

The volume group currently being edited can be modified in the selection box above to the left. The buttons above to the right enable you to create additional volume groups and to delete existing volume groups. However, only volume groups without any more partitions assigned to them can be removed. For a

standard SuSE Linux system that is installed, you do not need to create more than one volume group. A partition assigned to a volume group is also called a physical volume (often abbreviated to PV). To add a previously unassigned partition to the volume group you selected, first select the partition and then click on the button ‘Add volume’ below the selection list. This allows the name of the volume group to be entered next to the partition selected. You should assign all partitions to a volume group meant for LVM, otherwise the space on the partition will remain unused. Before you can exit the dialog, you will have had to assign at least one physical volume to each volume group.

3.8.4 Logical Volumes

Add, edit, or remove logical volumes in this dialog. Click ‘Add’ if you want to create a logical volume. Specify a size, format (reiserFS or ext2, for example), and a mountpoint in your file system for the volume.

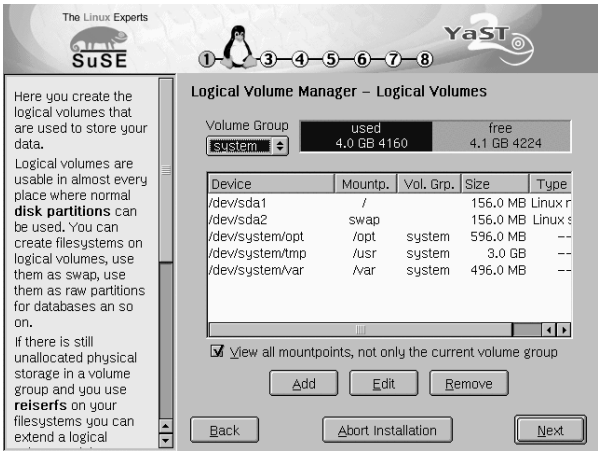



Figure 3.13: YaST2: Management of the Logical Volumes

If you have created several volume groups, you can switch between the different volume groups in the selection list above to the left. The new logical volumes are all located in the volume group shown at the upper left. After you have created all the logical volumes as they are required, the LVM configuration will be complete. You can then exit the dialog and continue on to software selection if you are currently in the process of installation.

Caution

Implementing the LVM is also associated with increased risk factors such as data loss. Possible dangers are application crashes, power outages, and faulty commands. Please secure your data before putting LVM to use, or before reconfiguring volumes – that is, do not work without backup!



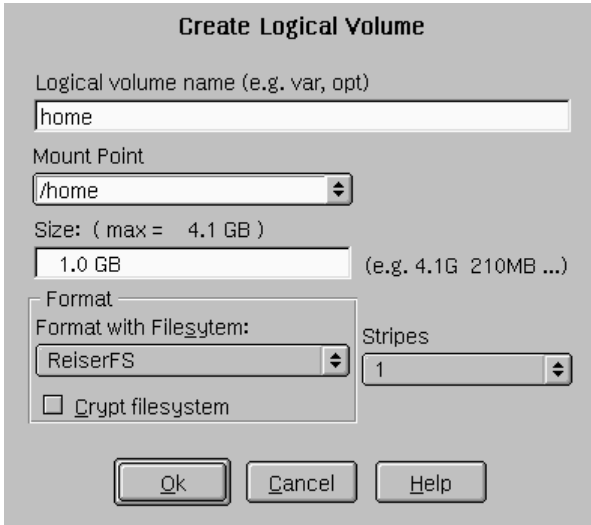


Figure 3.14: YaST2: Creating Logical Volumes

3.9 Configuring the Crypto File System


Partitioning in YaST2 gives you the option of completely encrypting a partition, that is, by creating a file system, which is subsequently encoded with the “twofish” algorithm. While the partition is being mounted, the data is not encrypted and can thus be read by anybody. Once it is unmounted, however, the data will be absolutely secure. Even if the hard disk or laptop get stolen, there is no way the data be retrieved without a password.

If you want to encrypt a partition, specify the beginning and ending cylinder, and finally, the desired partition size, as suggested in the field shown, in the dialog window for creating partitions (Figure 3.15 on the following page). The mountpoint, where the encrypted partition is to be accessed, can be arbitrarily defined. Now click on the item ‘File system encryption’ to the right and enter ‘OK’.

You will now be asked for the password in the next dialog window, which is confirmed by entering it twice. It must be at least five characters long and should be a combination of upper- and lower-case letters or numbers.

Caution

Be especially careful when entering the password here. This password cannot be changed later. If you forget it, your data will be irrevocably lost!!



Once this is accomplished, the new partition will now appear in the partitioning table, where the entry ‘CF’ for “Crypto Filesystem” will appear in the column marked ‘F’ (see Figure 3.16 on page 25).

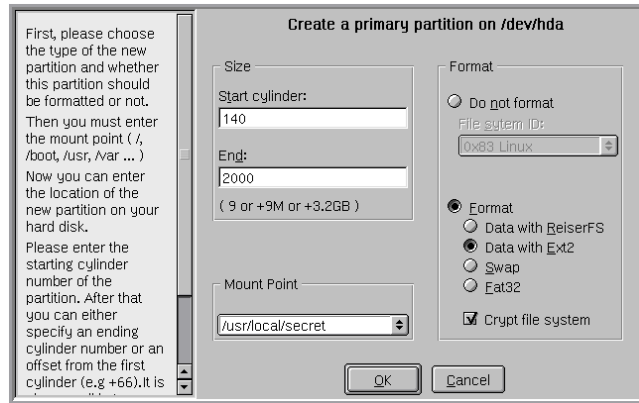


Figure 3.15: YaST2: Defining the Crypto File System

3.10 Boot Manager for System Start-Up

A boot mechanism is necessary for Linux to be able to start at all. The point in the system to which the boot manager LILO (Linux LOader) is to be installed must be defined here, as well as whether another boot concept should be applied. Normally, the SuSE Linux Connectivity Server is the only system installed on the machine and is also only installed on a single hard disk. If this is the case, it is best to install LILO to the boot sector (MBR) of the hard disk.

Otherwise, there is always the reliable method of creating a boot disk.

The menu 'Other configuration' also provides other options; see Figure 3.17 on page 26.

After installing, LILO can be reconfigured with the help of YaST2, or another boot floppy generated. More information can be found in Chapter 4.5.7 on page 61.

LILO: Other Boot Configuration

YaST2 now provides four options to select:

'To C: (in the MBR of the first hard disk)' – If SuSE Linux is to be installed as standalone operating system, LILO should definitely be installed to the MBR (Master Boot Record).

In the MBR, LILO can also act as boot manager for multiple operating systems. Only select this option if you are *certain* that your already installed systems are bootable from LILO— usually, Windows 95/98 has this capability. If you are in doubt, select the option 'Create boot floppy'.

'Create boot floppy' – If your machine is to run with multiple operating systems, generate a boot floppy for SuSE Linux. The previous boot mechanism is thus left unchanged. SuSE Linux can be booted from this floppy at any time.

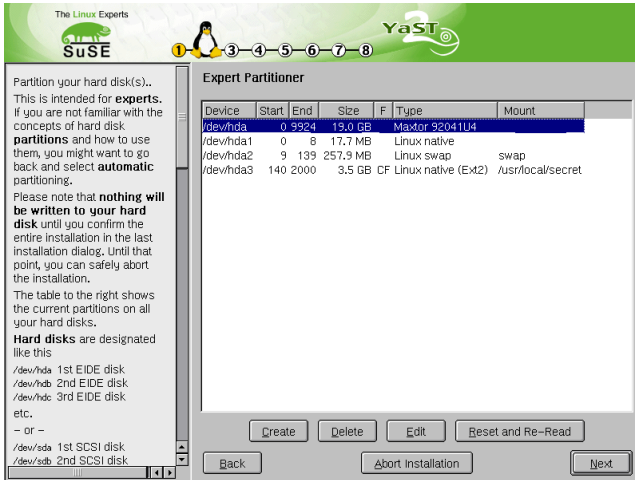


Figure 3.16: YaST2: Encrypted Partition

‘Do not install LILO (other boot manager)’ – Here, you can continue to use your own boot manager. Nothing is changed in the MBR (Master Boot Record); LILO will be configured on the `/boot` partition. However, in this case, you will be on your own in reconfiguring the existing boot manager.

‘To another partition’ – Select this option if you want to — or have to — specify another partition variant; see the previous item.

You will only need to fill out the remaining fields under specific circumstances. In case of doubt, please consult the YaST2 Online Help.

3.11 Root Password


root is the name of the superuser, or the system administrator. **root** is permitted to do all the things that the normal user is not permitted to do. The superuser may make changes to the system such as installing new applications or setting up new hardware. If someone has forgotten their password or has problems with software, **root** is able to help them.

For verification purposes, the password has to be entered twice (Figure 3.18 on the following page). Be particularly careful not to forget the root password.

Caution

If you forget the root password, it can be quite complicated to restore your system. Do not keep the password where a third party could have access to it.

Due to security reasons, we recommend that you do *not* log in as **root**. Use the administrator account for this purpose (see Section 3.12).



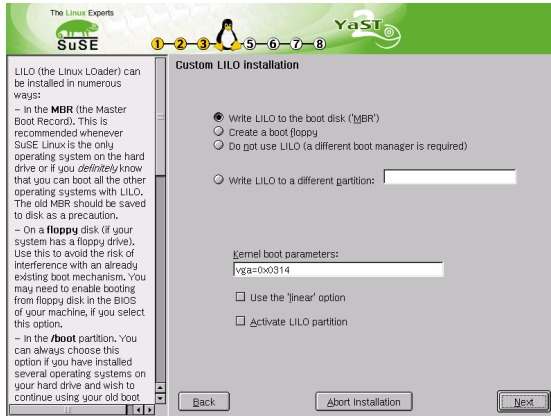


Figure 3.17: LILO Other Boot Configuration

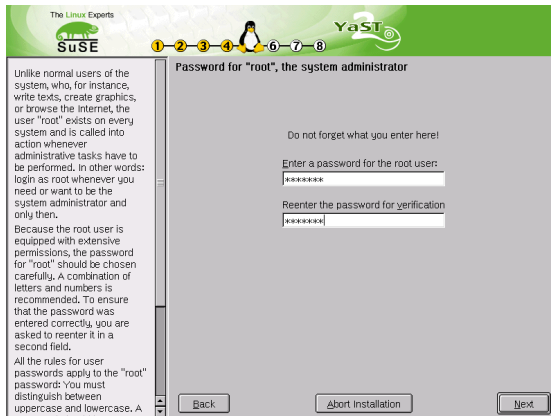


Figure 3.18: Setting the password for user root

3.12 Creating an Administrator Account

Once you have assigned a password for root, you will have to create your administrator account. This account is for taking care of daily tasks. Give yourself a memorable login name, which can be your first or last name, but not including any special characters or spaces. In conclusion, confirm your password by entering it twice.

Use a combination of lower- and upper-case letters as well as numbers for your password. Next, log into the system using this account.

In contrast to a normal user account, the administrator account provides certain features which simplify the administration of the SuSE Linux Connectivity Server.

3.13 Let's Go!

In the following dialog box (Figure 3.19), you will see your previously chosen settings listed. You can also 'Abort installation' here. The installation of SuSE Linux will then be ended and your system will remain unchanged. If you want to change some of your settings, you can click repeatedly on 'Back' until you've reached the dialog box where you want to make your changes. If you click on 'Next', however, a dialog box will appear, asking you if you are sure you want to proceed with the installation. If you answer 'Yes - install', the installation will begin. If you want to save your selections for later retrieval, click on 'Save settings to floppy' and all the installation settings will be saved to a disk.

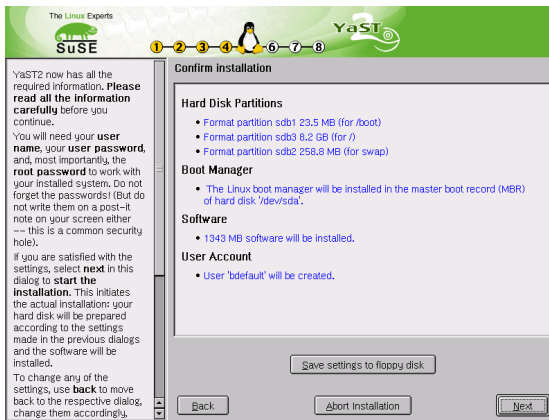



Figure 3.19: List of changes made

Caution

All data on the partitions you have specified for SuSE Linux will, in the next step, be deleted irrevocably. If you have chosen the entire hard disk, all other operating systems and data here will be erased.



3.14 Preparing the Hard Disk

YaST2 will now begin its work. YaST2 will create the selected partitions and format them. Depending on your system configuration, this may take some time.

3.15 Installation of Software Packages

Once you have started the installation process, the selected packages of the Linux base system are copied from CD or DVD and written to your hard disk. On this

screen, you can monitor the progress of the various tasks (Figure 3.20).

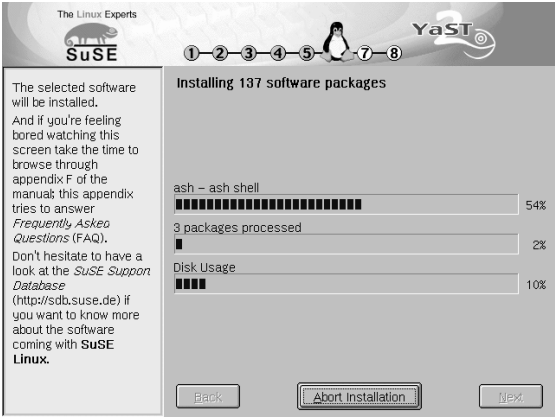


Figure 3.20: Package installation

Depending on the system configuration, the installation can be somewhat time-consuming.

To complete the installation of software packages, LILO is installed and a Linux base system started. Several messages will then appear on the screen.

Note

Depending on the configuration you selected for the installation of LILO, you might be prompted to insert a disk to create a boot floppy. Please note when doing this, that all data stored on the medium will be deleted.

3.16 Monitor Settings

If the installed monitor has not been automatically recognized, select the model from the list shown, refer to Figure 3.21 on the next page.

Some technical data regarding your selected model, the horizontal (HSync) and the vertical (VSync) frequency deflection rates will appear in the bottom portion of the screen. If the preferred model is not included in the list, you can manually enter the data in the entry fields, or choose pre-defined settings (VESA modes). Please use the relevant values listed in your monitor manual. Otherwise, you can use a driver floppy. To do this, click on 'Driver disk'. Insert the disk into the drive and confirm with 'OK'. If no file could be found or if the floppy is not readable, you will receive the respective warning. Then, the monitor data will appear in the selection list.

In the following display (Figure 3.22 on the facing page), define whether SuSE Linux should run in 'text mode' or in 'graphical mode' in the future. In the

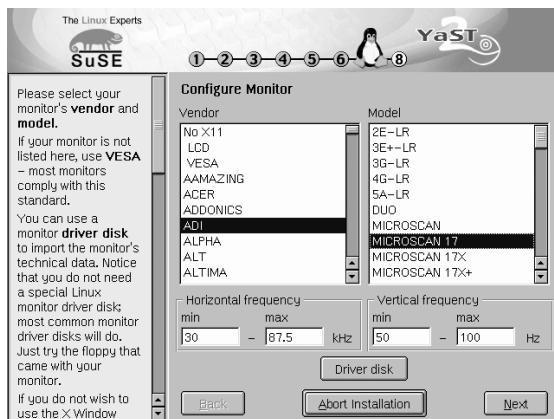


Figure 3.21: Selecting the monitor model

case of the SuSE Linux Connectivity Server, running it in a graphical interface is advisable for reasons of user-friendliness.

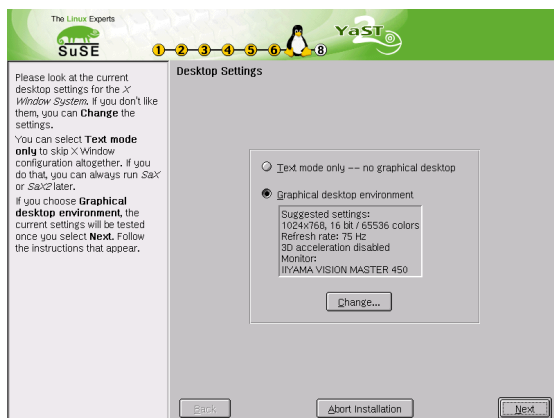


Figure 3.22: Monitor settings

By clicking ‘Change’, you will have the option of configuring the graphical interface (Figure 3.23 on the next page).

Note

If “3D acceleration” is listed, you must click on ‘Change’ in order to deactivate it, since this could lead to problems and is not required by the server anyway.



You can set the screen resolution and color depth for the graphical mode. You can even define the image repetition rate. By clicking on the ‘Test’ button, you

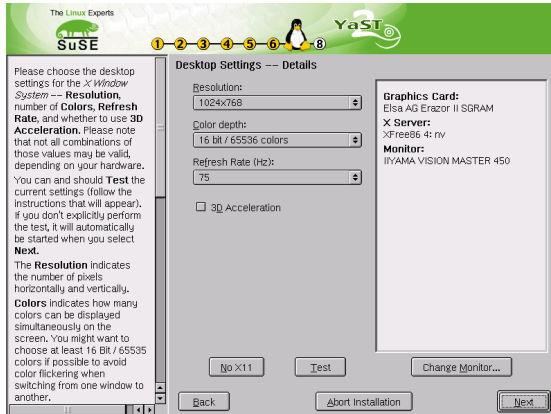


Figure 3.23: Changing the settings for the graphical interface

can test the resolution you have selected. The installation program will issue a message telling you that the screen will now switch over to the new resolution. If you don't see a steady screen, please stop the test immediately by pressing **(ESC)**.

3.17 Network Card

The first step in setting up the network is configuring the server's network card, which will be connected to the internal network. YaST2 automatically recognizes all network cards on your system and displays a list of these, as shown in Figure 3.24:

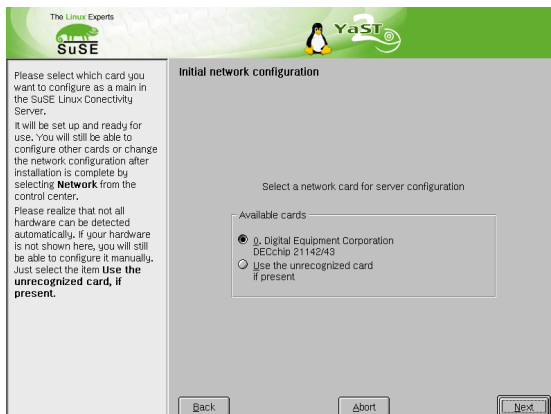


Figure 3.24: Selecting the network card

The card selected at this juncture is to be connected to your local network. All server services will only be available over this interface (=eth0, the first network card in Linux).

If you have more than one network card, they will usually be located in the upper slot on the computer, normally the outer one to the left on PCs.

If YaST2 did not automatically recognize the network card, it can still be configured manually by clicking 'Use non-recognized card, if it exists'. YaST2 allows you to enter the driver name by hand (or the kernel module), which is required by your network card. Clicking 'Select from list' will give you the option of choosing a driver from this list.



Figure 3.25: Network Card: Manual Configuration

Assigning a Host Name

The name which defines your computer in the network must be entered in this YaST2 screen. This name consists of the actual host name and the domain name. Any part of this name may contain letters, numbers and the '-' symbol. The domain name is made up of several components, separated by periods.

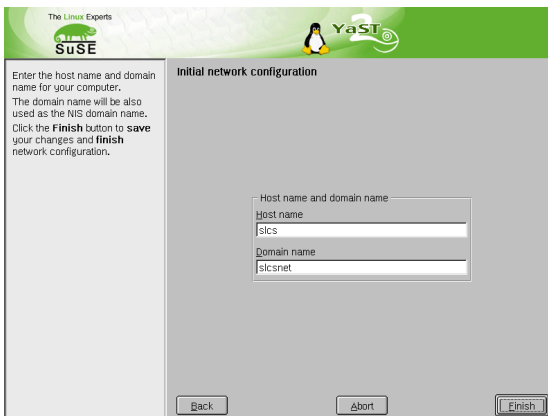


Figure 3.26: Host and Domain Name

The host name is the name the computer has in the network, such as `slcs`. The name should not contain any more than eight characters and may not be given more than once in a local network.

The domain selected here describes the local network and is predefined by the value `slcsnet`. It is wholly independent of NT domains as well as the Internet domain. The local domain is responsible for identifying the host when the TCP/IP protocol is being used and will automatically be forwarded to the clients connected to the local network.

NT domains (by default, set to `workgroup`) have a similar function, but only apply to the proprietary SMB protocol such as that used by Microsoft Windows.

The Internet domain serves to identify a network in the Internet and must therefore be registered. However, since this involves only a local network, which is not reachable from the outside thanks to the firewall, the Internet domain is irrelevant in this context.

3.18 Finishing the Installation

As soon as the SuSE Linux Connectivity Server basic configuration is completed, the Linux system will reboot to its final operational state, at which point, numerous messages will be reissued on the screen.

Once installation is finished, ‘`SuSEconfig`’ will run, in order to initialize the running SuSE Linux System.

Finally, you should definitely browse the ‘`Installation Protocol`’, to make sure all steps were completed successfully with ‘`OK`’.

Caution

Once you have completed installation, all server services will be activated (for more on this, see the tip on page 11)!



3.19 Graphical Login

SuSE Linux is now installed and configured so that you can log on to your system. Your monitor will now display the graphical login. Enter next to the username the login name of the administrator account you specified earlier in Section 3.12 on page 26.

Caution

Due to security reasons, we discourage starting the graphical interface as `root` (see Chapter 3.11 on page 25). We advise you to only log in as ‘`root`’ in absolutely dire circumstances.



If your login was successful, the desktop environment will be started. Your administrator account already provides several icons on your desktop.

If administration tasks are on hand, click on the 'YaST' icon. A window will open up where you must enter the root password. Once YaST2 has been started, you can carry out your configuration tasks.

If other users require use of your network, more explanations on this procedure can be found in [Chapter 4.5](#) on page [53](#).

4 SLCS Server Configuration with YaST2

With the help of YaST2, enhance your SuSE Linux system with additional hardware components, such as a printer or sound card, configure system services and the network, and install or remove software.

Many Paths to YaST2

Via the K Menu, there are several ways of accessing YaST2: via the 'Control Center', via 'SuSE' → 'Administration' → 'Configuration', and via 'Preferences'. Otherwise, change to user 'root' (**su** - then enter the root password) in the shell and enter **yast2**.

In the K menu pop-up menus, directly click on the configuration module needed. YaST2 will open a small dialog once it is loaded. Here, enter the password for user 'root' (the system administrator). The configuration then will be carried out as user 'root', because only root is permitted to make changes to the Linux system files.

Note



As a reminder: you should only be logged in on the computer as root for administrative tasks such as maintenance and system repairs. Being logged in as root is too risky for daily operation, since root can irrevocably delete all files.

If, for whatever reason, you are not able to run YaST2 as described above, there is a slightly more complicated way to do this. Enter the following in a shell on the graphical desktop:

```
xhost +
su -
(enter root password)
export DISPLAY=:0.0
yast2
```

After exiting YaST2 switch back to normal user from 'root' with **exit** and then enter **xhost -** to reactivate the access controls for the X server.

To change the language for YaST2 select 'System' then 'Choose language' in the YaST2 Control Center. Select the desired language then exit YaST2 and restart it.

The YaST2 Control Center

Next, the YaST2 Control Center will appear. The area to the left of the screen is divided into 'Hardware', 'Network/Basic', 'Network/Advanced', 'Security/Users', 'software', 'System', and 'Miscellaneous'. If you click one of the icons, the respective contents will be listed to the right. For example, click on 'Sound' and a window will open where you can make configurations for your sound card.

Configuration takes place in several steps. YaST2 guides you through all the dialogs with 'Next'. In the left portion of the screen, a help text is displayed regarding the respective topic, explaining the entries required. Once you have completed the necessary entries, use 'Finish' to complete the last configuration dialog. The configuration is then saved.

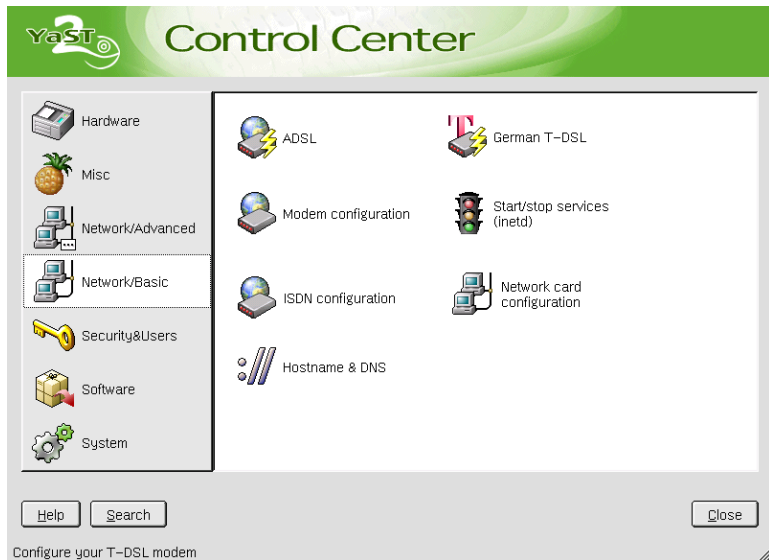


Figure 4.1: YaST2 System Configuration and Administration

4.1 Hardware

Before starting the software configuration for new hardware, the hardware itself needs to be installed. Follow the instructions provided by the vendor. Switch on external devices, such as printers or modems, and open the respective module in YaST2.

Most of the hardware is auto-detected by YaST2, so only a few additional settings have to be done manually to get the hardware running. If auto-detection fails, YaST2 provides a list of devices from which to select the appropriate device. Consult your hardware documentation if the information printed on the device itself is not sufficient.



Note

Beware of model descriptions. Try a similar description if you do not find your model in the device list.

In some cases, however, exact specifications to the number or letter are absolutely necessary, since more general descriptions cannot always guarantee compatibility. Unfortunately, even similar hardware does not always understand the same language.

4.1.1 Printer Configuration

Add and configure local and network printers with ease in YaST2. To do this, click on 'Printer' in the start screen. YaST2 will now load the necessary settings for printer configuration (Figure 4.2).

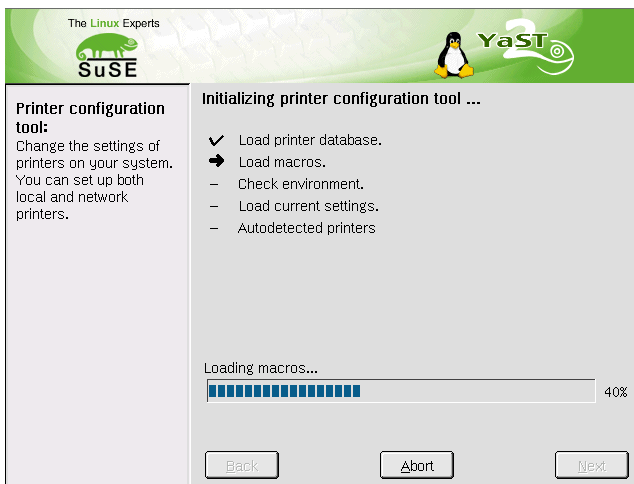


Figure 4.2: YaST2: Initializing the Printer Configuration Tool

Next, a list of active printers already connected to your computer, or to your network, will be shown. Now click on 'Add' and choose whether you want to install a local printer, a printer from the Linux network, or from another network (Novell or Samba) (Figure 4.3 on the next page). Select the desired category then click 'Next'.

If you want to configure a printer integrated into your network, for example, you must specify a print server. With a click on the double-arrow next to the text field, you will be presented with a list of available hosts and printer names. If you want to use a print server or network printer not included in the list, you will have to know its name and IP address. As soon as you have selected or specified one, you can select 'Test' to check to see whether it is, in fact, a printer or print

server, as well as whether it can be reached. If the print server was properly detected, YaST2 will prompt you in the following dialog to give a name. If, on the other hand, no print server was detected, an error message will appear instead.

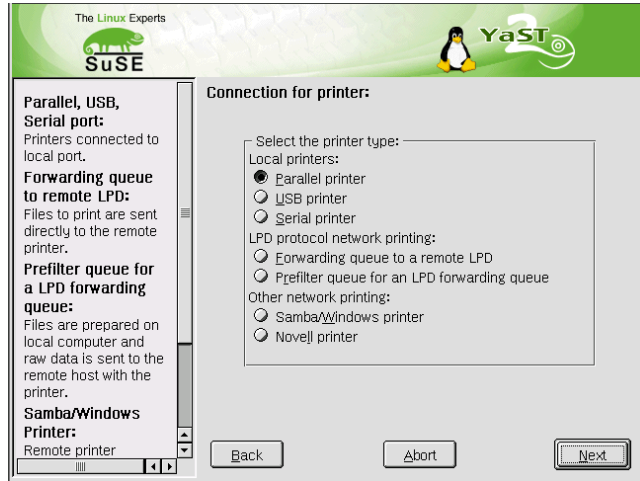


Figure 4.3: YaST2: Selecting the Printer Type

Integrating a printer from a Samba or a Novell network is similar. Once again, you will have to specify a print server or select one from the list. The difference to the Linux network is that, in this case, user information does not need to be known or specified.

If you want to connect a local printer to your parallel port, select the item 'Printer on Parallel Port' after clicking on 'Add' then click on 'Next'. Now choose the parallel port connection. With 'Test', you can again review the printer connection.

If the test was successful, a list of the most popular commercial printers will be displayed in the next window. Select your model. Information is available about the Linux support, depending on the printer model, and for GDI printers, information on where to obtain a Linux driver (see Chapter 4.4 on the facing page). Local printers can also be integrated into a serial or a USB interface the same way.

GDI Printer Issues

Many printers are sold as "Windows printers" or GDI printers. (GDI stands for the Windows Graphical Device Interface: such printers are designed to work with only one operating system.) They are often difficult or impossible to configure to work with Linux: some of them are capable of using other standard printer languages and are thus usable, while others will only work at all with Windows (TM). Consult the CDB at <http://cdb.suse.de/> or check with the hardware manufacturer if you are unsure.

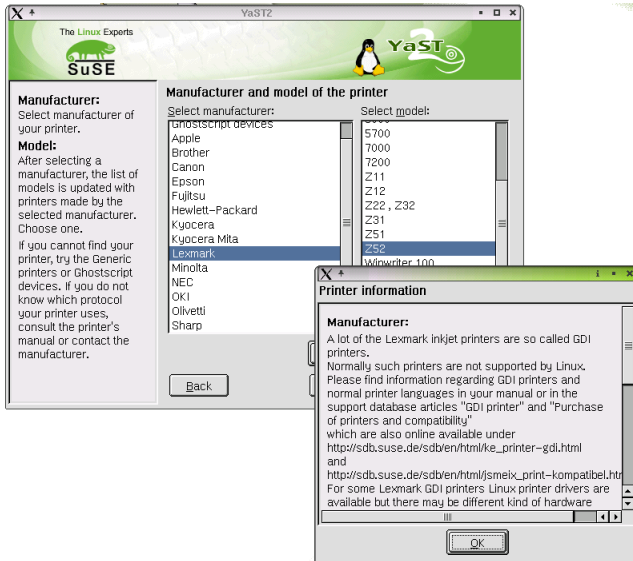


Figure 4.4: YaST2: Selection of the Local Printer with Infobox

With GDI printers, the manufacturer does without a standard protocol completely and controls the printer directly with control sequences of the specific model. However, there are printers on the market which can act both as GDI printers and also work with “proper” printer languages.

4.1.2 Graphical Interface (X11)

The graphical interface, the X11 System, provides the user with the basis for working in a graphical environment, as the graphical user environment (such as the KDE desktop) runs on top of the graphical interface. The X11 settings are saved in files which vary according to the XFree86 version being used:

```
XFree86 3.x: /etc/XF86Config
XFree86 4.x: /etc/X11/XF86Config
```

The graphical interface is usually configured during installation. However, if you still want to improve the values or connect another monitor to a running system, use this YaST2 module. The current configuration will be backed up before changes are made. The start screen will allow restoration of a saved previous X11 configuration. Then, you will be taken to the same dialog as in the SuSE Linux installation. You have the choice between text mode and the graphical interface. The current values will be shown for the graphical interface: the screen resolution, color depth, image repetition rate, vendor and monitor type (if this has been auto-detected), and, possibly, an already existing 3D acceleration. Click ‘Change’ to configure the monitor. If you have a graphics card with a 3D chip, enable 3D acceleration here. Depending on the hardware you are using,

when selecting the color depth, you will have the option of choosing 16, 256, 32768, 65536, and 16.7 million colors at 8, 16, or 24 bits. At least 256 colors is recommended.

Test the settings by clicking 'Test'. If you click 'Next' right afterwards, the test runs automatically. If you are not getting a still picture, stop the test immediately with **(Esc)** and reduce the values. Use the test image to adjust the dimensions and position of the screen display. Test it using the small white squares located in the four corners of the test screen. These should be fully visible without color distortions for an optimal screen position.

If your monitor is not automatically recognized, you will be taken to the monitor selection dialog. Also reach this dialog with 'Set monitor specifications'. The vendor and device list offers a large selection of models, where you will most likely find your monitor, manually enter the values for your monitor, or choose the default settings, VESA modes.



Caution

Be extremely careful with manually entering the permissible deflection frequencies. The wrong values could destroy your monitor. Look up the values in your monitor manual.

To be safe, choose a standard resolution to start. Highlight the item 'Vesa' and the values 640x480. The Vesa mode is, however, limited to a 75-Hz image repetition rate. For modern monitors, anywhere between 75 and 90 Hz is a suitable repetition rate. Sometimes, display errors can be attributed to hardware limitations. Alternatively, you may have the option of using the existing driver disk. To do this, click on 'Driver floppy', insert the monitor vendor's floppy, and confirm with 'OK'. If this works, the monitor data will then appear in the selection list.

4.1.3 Keyboard

The preferred keyboard layout usually corresponds to the selected language. Use the test field to try out the configuration. Make sure that the 'z', 'y', and special characters are correct on your keyboard.

4.2 Internet Access

4.2.1 Basic Internet Connection

All the machines on the Internet make up a large network where various operating systems are running with different hardware. The Internet uses a standard communication protocol that can be understood regardless of hardware or software used. This is done by the Internet Protocol (IP), together with the Transmission Control Protocol (TCP), the User Datagram Protocol (UDP), and the

Internet Control Message Protocol (ICMP). These protocols comprise the common “language” used by all machines on the Internet. The abbreviation for this is TCP/IP.

Every machine on the Internet has an ID number, the IP address. It can only be addressed by TCP/IP with this number. Normally, a machine also has a text name, used by application programs to refer to them. The Domain Name System (DNS) is responsible for converting the IP address to a text name. This particular service is offered by name servers. A machine or an application offering a service is called a server (for instance, DNS server) and a machine or application making use of a service is called a client.

Below TCP/IP, there are various standardized protocols for forwarding the appropriate TCP/IP data transfers to the given transmission method. For network connections via a network card, this is the ethernet protocol. For modem and ISDN telephone connections, it is the Point to Point Protocol (PPP), and for ADSL/T-DSL connections, the Point to Point over Ethernet Protocol (PPPoE).

The ethernet, PPP, or PPPoE connection, followed by the TCP/IP connection between your own machine and a machine on Internet provider, must be established before setting up an Internet connection.

On top of TCP/IP, there are various standardized protocols for proper data transfer to the application.

- The HyperText Transfer Protocol (HTTP) serves for the transfer of web sites in HyperText Markup Language (HTML) format.
- The Simple Mail Transfer Protocol (SMTP) is responsible for sending e-mails to another machine and Post Office Protocol (POP3) for downloading e-mails from a mail server.
- The File Transfer Protocol (FTP) is used to transfer files.

For several application programs, such as a web browser and an e-mail program, to use the same Internet connection at the same time, separate TCP/IP connections are used for each application. Large amounts of TCP/IP data are also split up into small packets, so that HTTP packets from the web browser can be sent over its TCP/IP connection while alternating with SMTP or POP3 packet transfers from the e-mail program via other TCP/IP connections.

Since several applications are using the same Internet connection, the IP address, which only identifies the machine, is not enough. A port number is needed to sort out which TCP/IP data belongs to which application.

These standard services are usually provided on their particular server at the following port numbers: DNS on port 53, HTTP on port 80, SMTP on port 25, POP3 on port 110, FTP on ports 20 and 21.

The client can only implement the right service if it addresses the correct port number at the server.

4.2.2 Instructions for all Types of Internet Access

Personal Firewall

The Personal Firewall is especially intended for preventing Internet machines from setting up a connection to your own machine, without much effort in the way of configuration. At the same time, however, connections originating from your own machine to hosts on the Internet are allowed. The Personal Firewall is well-suited and more than sufficient for meeting customary demands. Only the name of the network interface (ppp0, ippp0, eth0) can be configured in the file `/etc/rc.config.d/security.rc.config`, where, in particular, connection requests are denied. YaST2 will take care of this for you if you click on the item ‘Enable firewall’ in the corresponding dialogs.

The following will be filtered out by the personal firewall:

- All TCP connection requests. The security is based on the fact that the personal firewall will always block the first incoming TCP packet, prevents a proper TCP connection from being established. Those TCP packets which are not a part of an existing TCP connection and are not TCP connection requests will be discarded in any case.
- All UDP packets, except for those on port 53 from one of the configured name servers (normally only the provider’s name server, usually automatically configured when the Internet connection is set up; refer to “Internet Connection and Local Network” on page 42.
- Some of the less conventional ICMP packets.

All filter rules only apply to the configured interfaces, and nothing else. Some services can lead to “side effects”. Among these are IRC (CTCP), FTP (PORT mode, passive FTP, used by customary browsers, works), printer services, real audio, real video, cucme, napster, ICQ, and a few others.

Automatic Dial-Up (Dial on Demand)

If you click on ‘Dial on demand’ or ‘Automatic Dial-in’ in the YaST2 modules, the Internet connection will be made automatically when required, for example, when you enter an external URL in the browser. ‘Dial on demand’ is only recommended if you have a flat-rate Internet connection, as processes running in the background, such as frequent e-mail retrieval, require regular dialing into the Internet.

4.2.3 Internet Connection and Local Network

In every Internet connection, there is a normal TCP/IP connection between the local host and a host at the Internet provider. Normally, use the DNS of your ISP. The network is configured so that the connection to the Internet provider is used for all TCP/IP data not intended for the local host. This is normally correct,

because the local host does not usually function as a DNS server and does not have any other network connections, so all TCP/IP data is Internet-related.

There are usually no problems on the network with the TCP/IP connection to the Internet provider, if there is only one local host. An exception is if, for example, a firewall has been configured so that no data can be transferred at all.

4.2.4 ISDN

ISDN configuration can be found under 'Network/Basic'. If your ISDN card is successfully auto-detected, a dialog appears where you can make your 'Selection of ISDN protocol'. 'Euro-ISDN (EDSS1)' is the standard for this (refer to Scenarios 1 and 2a below) in Europe. '1TR6' is a protocol used by older and larger phone systems (refer to Scenario 2b below). 'NI1' is the standard in the USA. If this automatic recognition fails, choose the correct ISDN card (Figure 4.5 on page 45). Then specify the ISDN protocol and go on to 'Next'. In the screen which follows, specify your country and provider. The ones listed here are "Call-by-Call" providers. If you want to use a provider not included in this list, click 'New'. The 'ISP parameters' screen will appear where you can make all the necessary settings pertaining to your preferred provider. 'ISDN SyncPPP' is the standard 'ISDN type'. Specify the provider name for the 'Connection Name' then the provider's telephone number. In the case of an interposed PBX, you might need an additional number in front of the phone number itself to dial out (usually a zero or nine, but it is best to refer to the instructions for your PBX). The entire telephone number may not contain any separators, such as commas or blank spaces. Enter the username and password received from your provider.

Next, proceed to the ISDN connection parameters. The following scenarios require various specification for your 'Phone Number':

1. The ISDN card is connected directly to the phone company's socket. Enter an "MSN," Multiple Subscriber Number, if provided by your phone company. Otherwise, leave it blank and the ISDN card should work.
2. The ISDN card is connected to a PBX:
 - a) The telephone system's protocol is Euro-ISDN/EDSS1 (usually for "small" phone systems for households): These phone systems have an internal S0 bus and use internal numbers for the connected devices. In this case, specify the internal number as MSN. Further information can be obtained from your phone system documentation. One of the MSNs available for your phone system should work as long as this MSN is allowed external access. If all else fails, a single zero might work as well.
 - b) The phone system's protocol for the internal ports is 1TR6 (mostly the case for "large" corporate telephone systems): the MSN is known here as "EAZ" and is usually the extension. Usually, you only need to enter the last digit of the EAZ for the Linux configuration. If all else fails, try the digits 1, 2, 3, 4, 5, 6, 7, 8, or 9.

Choose a dial mode as follows: 'Manual', 'Automatic', or 'Off'. Look at page 42 regarding the 'Automatic' dial mode. It is best to choose 'Manual',

because, afterwards, you can conveniently dial into the Internet using kinternet, for example. Dial in a command line with

```
/usr/sbin/isdnctrl dial ippp0
```

and hang up with

```
/usr/sbin/isdnctrl hangup ippp0
```



Note

Be careful with the 'Automatic' dialing mode unless you have a flat-rate connection.

You can also configure after how many seconds the connection should be terminated if data transfer is no longer taking place. 60 seconds are recommended for this. Along these lines, when enabled, 'ChargeHUP' also exists to make sure that the connection is not terminated until the next payable unit. However, this does not work with every provider.

It is highly recommended to select the item 'Initialize ISDN System when booting' so that the necessary drivers are loaded. This alone will not set up an Internet connection. You can also enable the firewall. This way, your machine will refuse external connection requests, while you can continue to use the network as normal. Note that there are two different firewall packages: the SuSE-firewall and the Personal Firewall. Unlike the SuSEfirewall, Personal Firewall cannot be custom-configured. The only specification which can be made for Personal Firewall is the name of the network interface (ippp0, eth0 etc.), where incoming packets can be blocked.

You should just accept the addresses suggested by YaST2 under 'IP Settings'. The preselected items 'Dynamic IP Address' and 'Dynamic DNS' ensure that the IP address and name server assigned by the provider are forwarded during the connection, which is usually necessary. Under 'Callback settings', 'Callback off' should be selected, as the other choices are — at least for personal computers — irrelevant.

'Next' and 'Finish' complete the configuration.

4.2.5 Modem

Normally, these days, companies no longer have dialup connections to the Internet over a modem, but rather, over DSL, ISDN, or a leased line. There are still ways to implement a modem dialup connection (Figure 4.6 on the facing page), but its configuration will only be briefly touched on here. Its configuration is — for the most part — intuitive, and is carried out in much the same way as the configuration for ISDN, as described in Section 4.2.4 on the page before. Settings for the modem can be made in respect to Baud rate and initialization strings in the 'Details' menu, where changes can be made, provided that you are familiar enough with what you are doing. However, it is generally not necessary to do this. You should only make changes in this menu if your modem was not auto-detected, in which case, it would have to be specially configured

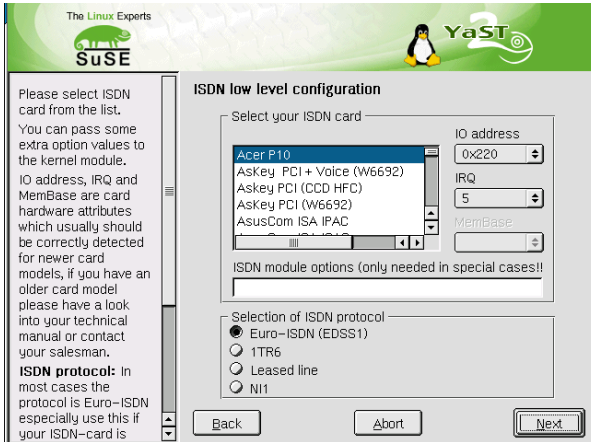


Figure 4.5: YaST2: ISDN Configuration

for data transfer. This is usually the case for what are known as “ISDN terminal adapters”.

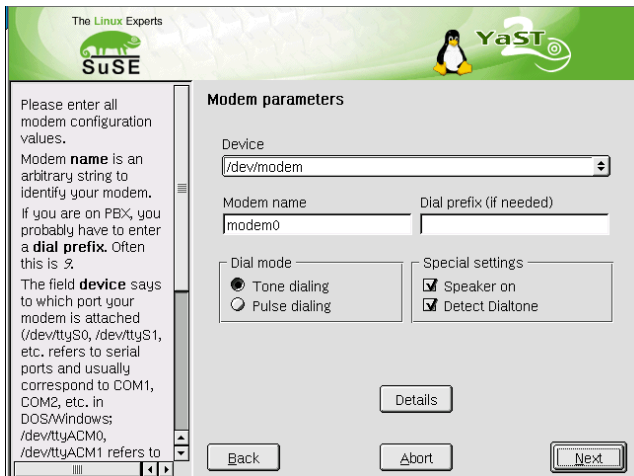


Figure 4.6: YaST2 Modem Configuration

4.3 Network Card

With the help of YaST2 you can configure additional network cards under ‘Netzwerk/Basic’ following installation.

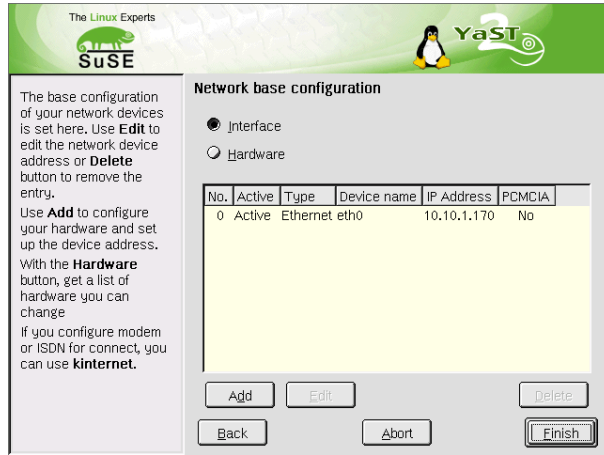


Figure 4.7: Dialog for Network Base Configuration

Caution

Only change the configuration of the interface `eth0` or the server's IP address if you know exactly what you are doing. The IP address must be left at `192.168.0.1` in order to correspond accurately with the other network service configurations.

The dialog shown in Figure 4.7 will appear. With 'Add', add the network card to the configuration. With 'Remove', remove it from the configuration. With 'Edit', modify the network card configuration.

Activate the item 'Hardware' to modify the hardware data for an already configured network card with 'Edit'. You will arrive at the menu for changing the settings of the network card. This menu is shown in Figure 4.8 on the next page.

Normally, the correct driver for your network card has already been configured by YaST2 during installation and is activated. Therefore, manual hardware parameter settings are only needed if you are using more than one network card or if the network hardware is not automatically recognized. In this case, select the item 'New' to specify a new driver module.

In this dialog, set the network card type and, if you have an ISA card, the interrupt to implement and the IO address. For some network drivers, you can also specify special parameters such as the interface to use or whether you to have an **RJ-45** or a **BNC** connection on your card. For this, refer to the driver module documentation.

After entering the hardware parameters, configure additional network interface data. Select the item 'Interface' in the dialog 'Network base configuration' to activate the network card just set up and assign it an IP address. Select the card number then click on 'Edit'. A new dialog will appear where you can

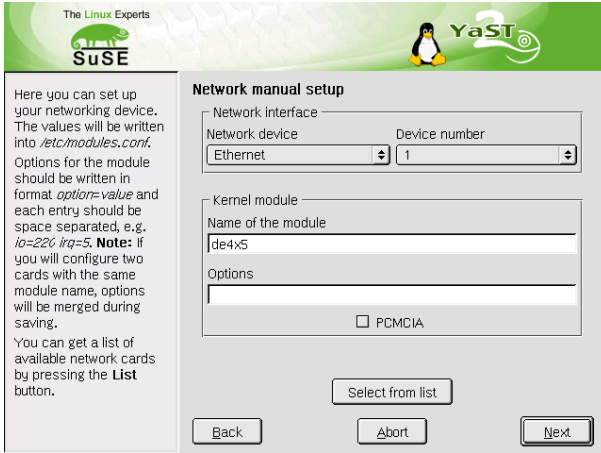


Figure 4.8: Configuration of the Hardware Parameters

specify the IP address and other IP network data. Look at Figure 4.9 for more information.

Select the card number then click 'Edit'. A new dialog will appear where the IP address and the rest of the IP network data can be specified. Normally, no additional information new needs to be entered here.

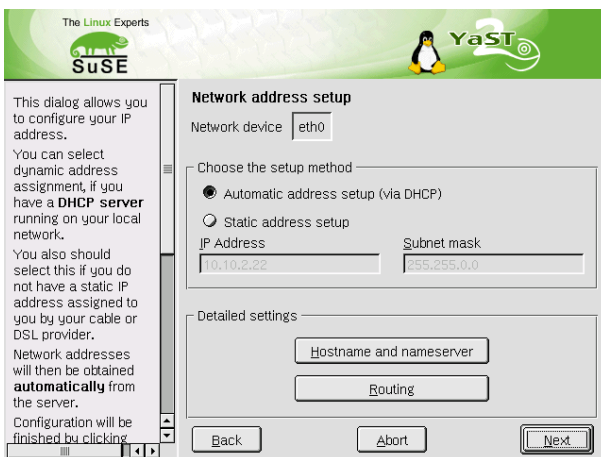


Figure 4.9: Configuration of Network Addresses

4.4 Server Services

4.4.1 Basic Samba Configuration

With the program package Samba, SuSE Linux Connectivity Server can be enhanced to a powerful file and print server for DOS and Windows machines as well.

Normally, no changes need to be made here, provided you have adopted the default values (domain `workgroup`). Once it is installed, SuSE Linux Connectivity Server will run as a primary NT domain server, or primary domain controller (PDC for short) under the domain name `workgroup`.

If you want to change the values anyway, first the Samba server's basic configuration must be defined. Here, specify whether the server should function as a work group server, or as (primary) domain controller.

Next, enter the appropriate name for the work group or domain. The description string serves to make the identification of the server easier when browsing through the network.

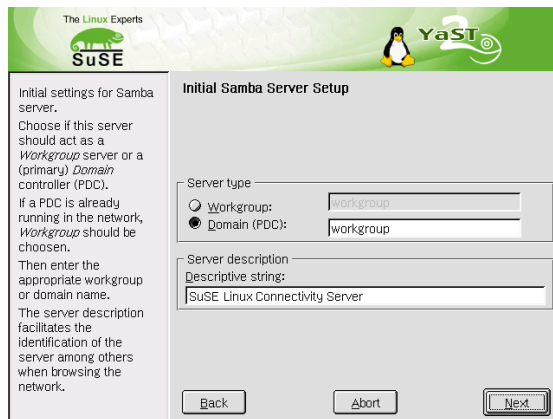


Figure 4.10: Samba Basic Configuration

Section 6.2.2 on page 87 provides more details. Further information can be found in the Samba book [BD00].

4.4.2 NFS Server Configuration

YaST2 enables you to quickly turn any host on your network into an NFS server. This is a server which makes the directories and files of all the hosts available to those permitted access to it. There are many applications, which can be provided for your employees, for example, without having to install them locally on their hosts.

For installation, select 'Network/Advanced' in YaST2 then 'NFS Server' (Figure 4.11 on the facing page).

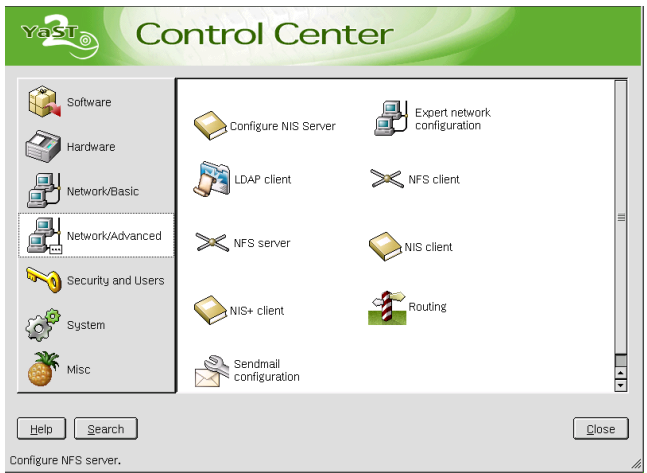


Figure 4.11: YaST2: NFS Server Configuration Tool

Next, activate ‘Start NFS Server’ and click on ‘Next’.

Now, only one step remains to be taken. In the upper text field, you will need to enter the directories to be exported. Then below, enter the hosts which are to have access to them (Figure 4.12). There are four options which can be set for each host: <single host>, <netgroups>, <wildcards> and <IP networks>. A more thorough explanation of these options is provided by the manpages on package exports (**man exports**).

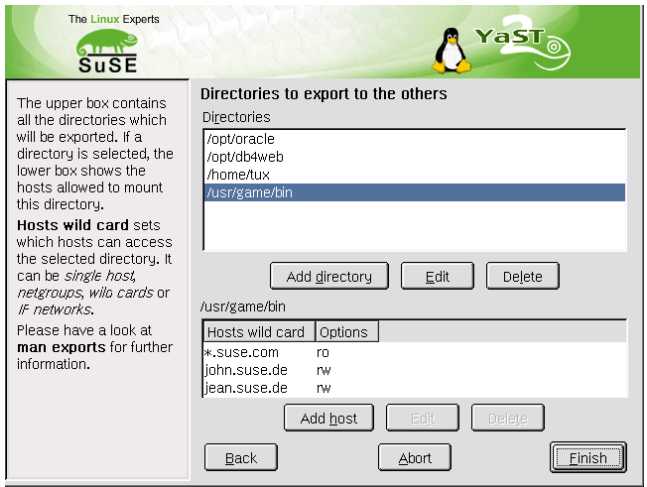


Figure 4.12: YaST2: NFS Server: Enter Export Directories and Hosts

‘Exit’ completes the configuration.

4.4.3 NIS – Network Information Service

What is NIS?

As soon as several Unix systems are to access shared resources in a network, user and group data has to be synchronous between hosts. The network should be transparent for the user. Regardless of which host he is working on, the user will always encounter the same environment. This is possible due to the services NIS and NFS. NFS is responsible for distributing file systems in the network. It was already described above, in Section 4.4.2 on page 48.

NIS (Network Information Service) can be described as a database service, enabling access to information from the `/etc/passwd`, `/etc/shadow` and `/etc/group` files network-wide. NIS can also be implemented for additional tasks as well (such as for `/etc/hosts` or `/etc/services`). However, this will not be discussed in detail here. A common term for NIS is 'YP', which is derived from *yellow pages*, meaning, the yellow pages of the network.

YaST2 NIS configuration

For installation, select 'Network/Advanced' in YaST2 then 'Configure NIS server'.

If an NIS server still does not exist on your network, you will first have to activate the item 'Configure NIS Master Server' in the next screen. If you already have an NIS server (that is, a "master"), add an NIS slave server if you are configuring a new subnetwork. First, you will be presented with an explanation of how to configure the master server. Enter the domain name at the top of the next configuration screen (Figure 4.13). In the checkbox underneath, define whether the host will also be an NIS client, that is, whether users, who can also access the data from the NIS server, will also be able to log in to it.

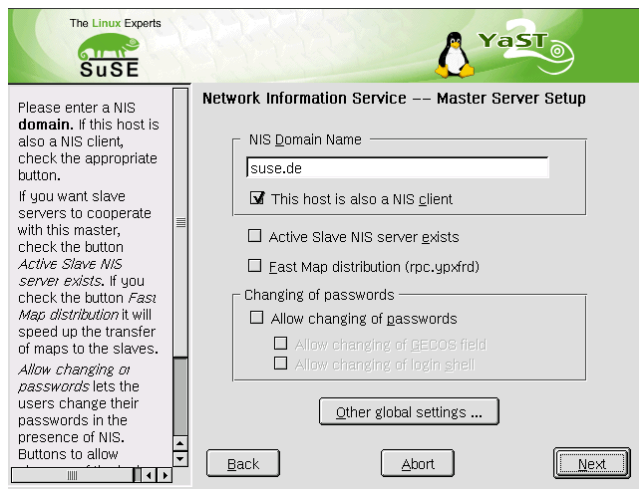


Figure 4.13: YaST2: NIS Server Configuration Tool

If, at a later point, you want to configure more NIS servers (“slave servers”) in your network, the box ‘Active NIS Slave Server exists’ must be activated. In addition, ‘Fast Map Distribution’ should also be activated, which will speed up the data transfer from the master to the slave server.

If you want to allow users in your network to be able to change their passwords (with the command `yppasswd`, that is, not just to change their local passwords, but also, those stored on the NIS server), enable this option as well, thereby activating the checkboxes ‘Allow modification of GECOS entries’ and ‘Allow modification of the SHELL entry’. “GECOS” means that the user can also change his name and address settings (with the command `ypchfn`). “SHELL” also means that the user can modify his default shell (with the command `ypchsh`, for instance, from `bash` to `sh`).

Under ‘Other global settings...’ a menu will appear (Figure 4.14), where the default directory (`/etc`) can be changed. In addition, passwords and groups can be consolidated here. The setting should be left at ‘Yes’ so that the files (`/etc/passwd` and `/etc/shadow` as well as `/etc/group` and `/etc/gshadow`) can be synchronized. Furthermore, the smallest user and group number can set. ‘OK’ returns you to the previous screen. Now click on ‘Next’.

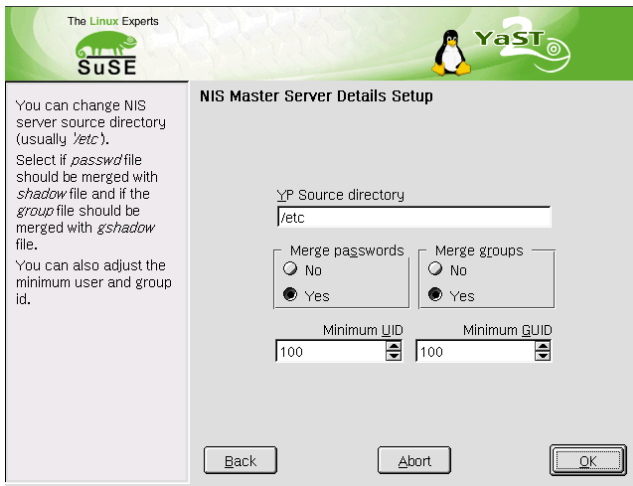


Figure 4.14: YaST2: NIS server: Changing the directory and synchronizing files

If you previously enabled ‘Active NIS Slave Server exists’, you must now give the host names to be used as slaves. Specify the name and go to ‘Next’. The menu that follows can be directly accessed, provided you did not activate the slave server setting before. Now the “maps”, the partial databases to be transferred from the NIS server to the individual clients, can be configured. The default settings can be applied under most circumstances, so nothing usually needs to be changed here. If you still want to make changes here, however, you should be very familiar with the material.

‘Next’ brings you to the last dialog, where you can define which networks are to be allowed to send requests to the NIS server (see Figure 4.15 on the next page).

Normally, this is your company network. If this is the case, there should be two entries:

```
255.0.0.0      127.0.0.0
0.0.0.0        0.0.0.0
```

The first one enables connections to your own host, while the second one allows all hosts, which have access to your network, to send requests to the server.

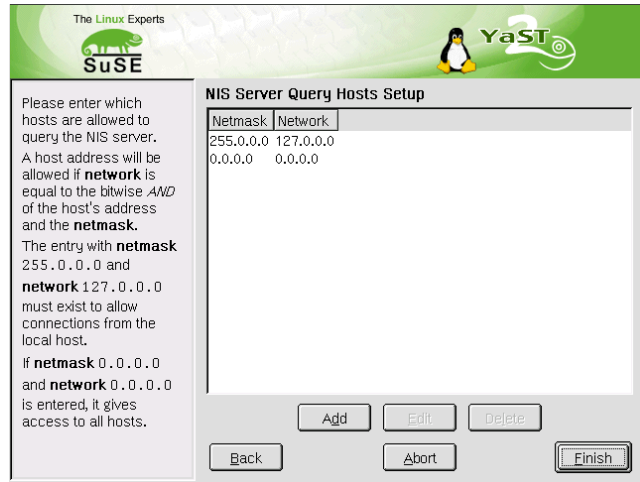


Figure 4.15: YaST2: NIS server: setting request permissions

4.4.4 E-mail (Sendmail)

In the configuration dialog located under 'Network/Advanced', the following items will be listed. Select the right one for you.

- 'Host with permanent network connection (SMTP)'
This is normally a "leased line", as is often found at companies or other institutions which work with the Internet. The Internet connection is always running so no dial-up is necessary. This menu item is also meant for members of a local network where no permanent Internet connection exists, but where a central mail server is used for sending e-mail.
- 'Single user machine without network connection'
If you do not have an Internet connection and do not belong to a network, you can only send e-mails locally.
- 'Host with temporary network connection (Modem or ISDN)'
Most home users need this option. It is for computers, not on a local network, that connect to the Internet via modem, T-DSL, ADSL, or ISDN.
- 'Use UUCP to send mail'
"UUCP" means "Unix to Unix Copy Program". In the past, it was often used for sending e-mails. This protocol is for dial-up connections and is not used as much these days.

- ‘Expert mode for sendmail configuration’
Proceeds to a custom configuration screen for expert settings, with ‘Next’.
- ‘Do not install /etc/sendmail.cf’
Select this item if a configuration already exists and it should not be changed.

The file `/etc/rc.config.d/sendmail.rc.config` is key for configuring sendmail. YaST2 configures this automatically according to the items selected. You can only (indirectly) access the contents of this file in the expert mode, where you can make changes to it by hand. The file `/etc/sendmail.cf` is generated with the help of a script read by sendmail. Exit the configuration with ‘Finish’.

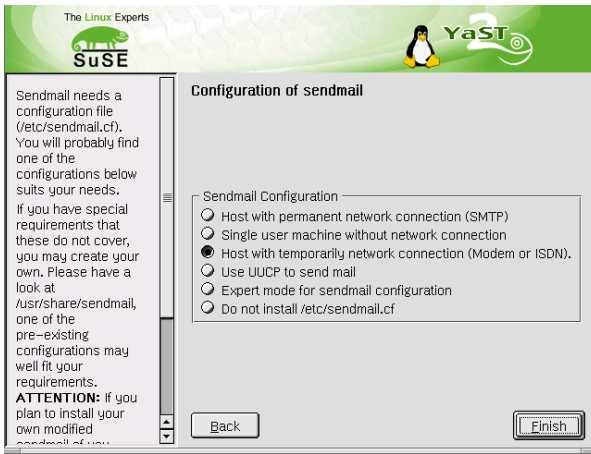


Figure 4.16: YaST2 Sendmail Configuration

4.5 System

4.5.1 Managing Users and Groups

Creating New Users

A basic aspect of Linux is that it is a multiuser system. Consequently, several users can work independently of one another on the same Linux system. Each user has a “user account” consisting of a user and login name and a personal password for logging in to the system. All users have their own home directories where personal files and configurations are stored.

In this module, located under ‘Security&Users’, easily add new users by simply filling out the fields as indicated then clicking ‘Add’. New users can log in to the system using their own login names and passwords.

‘Details’ offers several options for specialized settings, which should be left alone if you are not familiar with it. Find a selection list of default groups, the

home directory path which can be changed, the user ID, and a list of login shells. Define additional group affiliations below. If a new user is should access to the modem, “dialout” and “uucp” (unix to unix copy program) has to be entered.

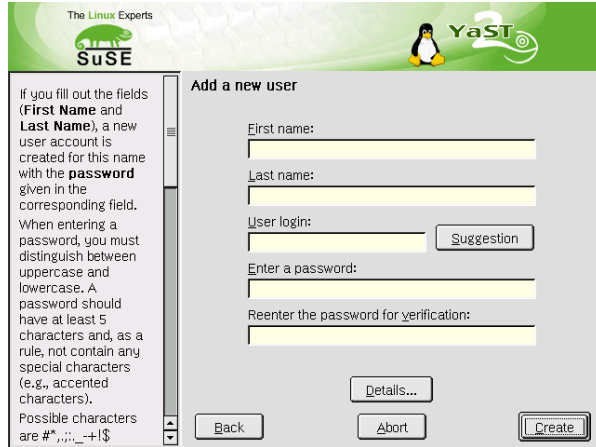


Figure 4.17: Adding new users with YaST2

Adding and Changing Users

After calling up this configuration tool, a screen will open labeled “Managing users and groups”. You will then be able to change users and groups. Group administration is under the “Changing and Adding Groups” module and is described there.

YaST2 provides a list of all users to assist in user administration (see Figure 4.18 on the facing page). To remove a user, simply click on the user in the list, so that the line is highlighted dark blue then click ‘Delete’. To ‘Add’ a user, proceed as described in “Adding New Users”. Under ‘Edit’, find the editing options under ‘Details’.

Creating a new group

Adding a new group is easy with YaST2 (see Figure 4.19 on page 56). For more information, read the YaST2 help text. When you specify members of a new group in the field below, be sure not to add any blank spaces before the commas separating the user and login names. YaST2 will suggest a group ID, which you can just accept.

Changing and adding groups

After opening this module, a screen will open: “User and Group Management”. You will then be presented with the option of either editing users or groups.

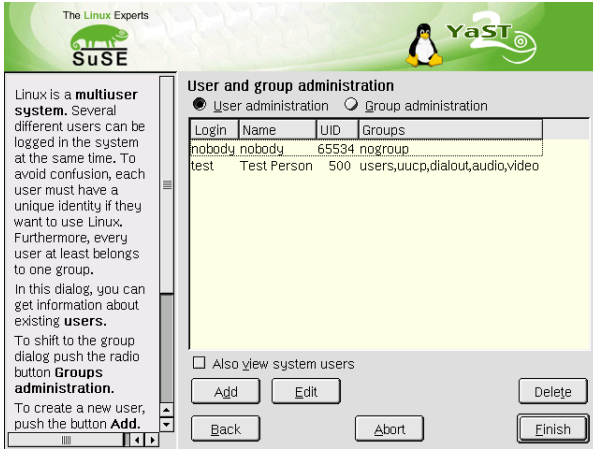


Figure 4.18: User Administration with YaST2

User administration is under the “Changing and adding users” module and was described there.

YaST2 offers a list of all groups to assist in group administration. To remove a group, click the group in the list, so that the line is highlighted dark blue, then click on ‘Delete’. To ‘Add’ or ‘Edit’ a group, proceed as directed in the YaST2 help texts displayed in the left pane.

4.5.2 System Security

In the start screen ‘Local security configuration’, which can be accessed under ‘Security&Users’, there are four selection items:

Level 1 is for stand-alone computers (preconfigured), Level 2 is for workstations with a network (preconfigured), Level 3 for server with a network (preconfigured), and custom defined is for your own settings.

If you click one of the three items, you will have the option of incorporating one of the levels of preconfigured system security options. To do this, simply click ‘Finish’. Under ‘Details’, access the individual settings which can be modified. If you choose ‘Custom settings’, you will be taken to the different dialogs with ‘Next’ automatically. Here, find the default installation values.

1. ‘Password settings’

Define how long the password should be for future users (minimum and maximum length). Five to eight characters is a reasonable number. Set for how long a password should be valid, when it expires, and how many days in advance an expiration warning should be issued (the warning is issued when logging into the text console).

2. ‘Boot settings’

This screen involves two things. First: How should the key combination



Figure 4.19: Adding a New Group with YaST2

(Ctrl) + (Alt) + (Del) be interpreted? Usually, this combination, entered in the text console, causes the system to restart. Leave it at that unless your machine or server is publicly accessible and you are afraid that someone could carry out this action without authorization. If you select 'Stop', this key combination will cause the system to shut down. With 'Ignore', this key combination will lose its affect entirely. Secondly: Who is permitted to shut down the system from KDM (KDE Display Manager — the graphical login)? 'Only root' (the system administrator), 'All users', 'Nobody', or 'Local users'? If 'Nobody' is selected, the system can only be shut down via the text console.

3. 'Login'

Typically, following a failed login attempt, there is a waiting period lasting a few seconds before another login is possible. The purpose of this is to make it more difficult for "password sniffers". In addition, you will have the option of activating the items 'Record failed login attempts' and 'Record successful login attempts'. If you suspect someone is trying to find out your password, check the entries in the system log files in `/var/log`.
4. 'Add user settings'

Every user has a numerical as well as an alphabetical user id. The correlation between these is established via the file `/etc/passwd` and should be as unique as possible. Using the data in this screen, define the range of numbers assigned to the numerical part of the user ID when a new user is added. A minimum of 500 is reasonable for users and should not fall short of this.
5. 'Miscellaneous settings'

For 'Setting of file permissions', there are three selection options: 'Easy', 'Secure', and 'Paranoid'. The first one should be sufficient for most users. The YaST2 help text will provide information on the three security levels. The 'Paranoid' setting is extremely restrictive and should serve

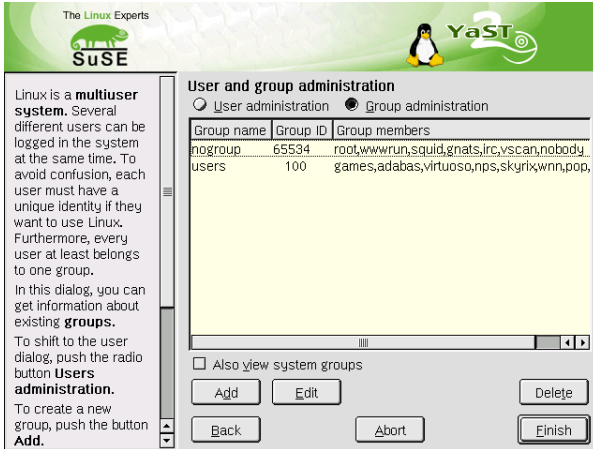


Figure 4.20: Group Administration with YaST2

as the basic level of operation for system administrator settings. If you select ‘Paranoid’, take into account possible disturbances and malfunctions when using certain programs, because you will no longer have the permissions to access various files. Also, in this dialog, define which users can start the “updatedb” program. This program, which automatically runs either on a daily basis or after booting, generates a database (locatedb) where the location of each file on your computer is stored (locatedb can be searched by running the **locate** command). If you select ‘Nobody’, any user can find only the paths in the database which can be seen by any other (unprivileged) user. If ‘root’ is selected, all local files will be indexed, since the user ‘root’, as superuser, may access all directories.

Another option is to activate the item ‘Omit current directory from the path of user root’, a reasonable selection. Finally, there is the option ‘Disable telnet login for user root’. It is also a good idea to choose this item. If not, any user on the network can log in to your machine as ‘root’ via telnet, through which the root password is deciphered to plain text.

With ‘Finish’, this configuration is complete.

4.5.3 Install and Remove Software

This module enables you to install more software on your machine. In addition, unwanted programs can be removed. To install from a CD, insert the first CD into the drive.

In the dialog, the package series will be shown to the left (commercial packages are often located in the ‘pay’ series). On the right, all the packages belonging to the series selected are listed. Packages already installed on your computer are marked with ‘i’.

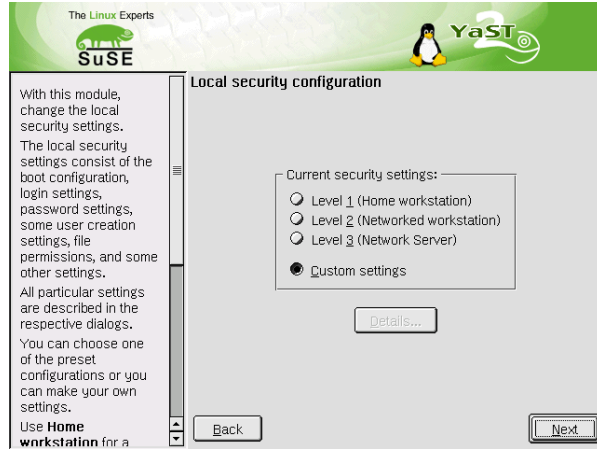


Figure 4.21: YaST2: System Security Configuration

Select and deselect a package by double-clicking or by selecting the line then clicking 'OK'. The packages selected for installation are marked with 'x' and ones to remove with 'd'. If a package requires additional packages, these will be automatically selected by YaST2 (label 'a') or you have the option of selecting one of several possible packages.

YaST2 evaluates the memory needed each time you choose an additional package. If the disk space is not sufficient, you will be informed by a warning window and one or more packages will have to be deselected.

If you exit the dialog with 'Cancel', your selection will not be saved and no actions will be carried out. With 'OK', the installation or removal of packages will be initiated. In the installation window, see the actions taking place via the progress bar. Once all packages have been processed, the installation will be completed by SuSEconfig. This can take some time. The hard disk normally becomes very active at this point.

Caution

You have the option of marking installed packages to be removed (these will be labeled with 'd'). Be aware of the warning messages while you are doing this. Do not remove any packages belonging to the Linux base system (series 'a').



4.5.4 Change Installation Source

The installation source is the medium on which the software to be installed is made available. Install from a CD (the usual route), from a network server, or from the hard disk. Read about this in the extensive YaST2 help text. When you exit the module with 'Save and exit', the settings will be saved and will be

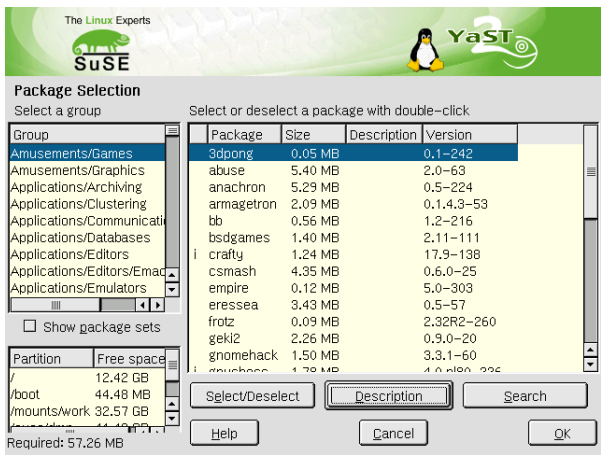


Figure 4.22: YaST2: Installing and Removing Software

applied to the configuration modules ‘Install/Remove Software’, ‘System Update’, and ‘Boot and kernel configuration’. These modules provide the option of continuing with ‘Install’ to install additional packages later or to remove them.

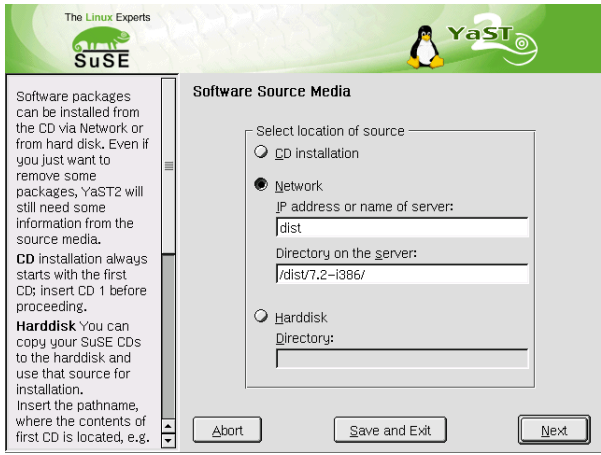


Figure 4.23: YaST2: Changing the Installation Source

4.5.5 Online Update

The YaST Online Update enables installation of important upgrades and improvements (see Figure 4.24 on the following page). Note that an online update can only be carried out if you have completed registration. You can find out more information on this in Section 2.1 on page 7.

The corresponding “patches” are available on the SuSE support server for downloading. The current packages can be installed automatically. On the other hand, you also have the option of personally specifying which patches to add to your SuSE Linux system via ‘Manual update’.

Click ‘Details’ to obtain information about your last update and the available packages. Find out about their contents by clicking on ‘Display patch information’. With ‘Next’, reach a list of all the available patches (if you chose ‘Manual update’), from which to make your selection. With ‘OK’ or by double-clicking, activate the individual objects. By clicking on ‘Next’ or ‘Finish’, the Online Update will be completed.

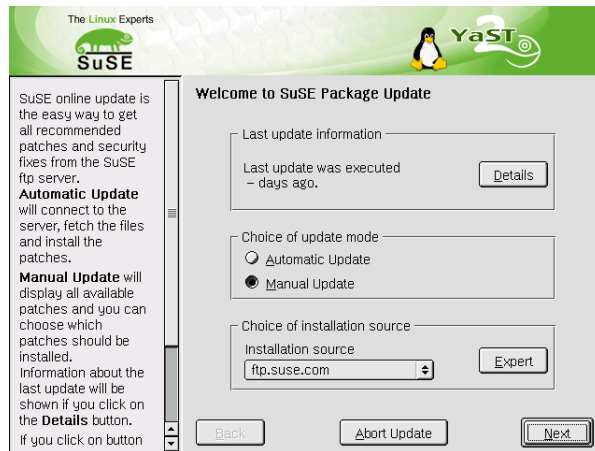


Figure 4.24: YaST2: Online Update

Online Update from the Console

To the benefit of system administrators and command line fans, the Online Update can be started in a shell. As ‘root’, load the current patch list and all related rpms from the first server in the `/etc/suseservers` list using the command:

```
earth:/root # yast2 online_update .auto.get
```

If you just want to load certain patches, you can add options to the command. Among these options are `security`, `recommended`, `document`, `YaST2` and `optional`. `security` retrieves security-related patches, `recommended` fetches updates recommended by SuSE, `document` provides you with information on the patches, or on the FTP server, `YaST2` fetches YaST2 patches, and `optional` gets minor updates. The command for downloading the security patches, for example, reads

```
earth:/root # yast2 online_update .auto.get security
```

The FTP server list from `/etc/suseservers` is typically loaded when you enter `.auto.get`. To disable it, deactivate the function in the `/etc/rc.config`. To do this, set `yes` to `no` in the line

```
earth:/root # YAST2_LOADFTPSEVER="yes"
```

The patches can now be installed with

```
earth:/root # yast2 online_update .auto.install
```

This command installs all fetched patches. To just install a group, use the same options as in `.auto.get`.

This method can be fully automated. The system administrator is able to download the packages overnight, for example, and then install the ones he needs the next morning.

4.5.6 System Update

Use this module to update and improve your system. It can be started at different stages in the process. YaST2 recognizes which packages need to be updated or you can decide on your own which package should be updated. However, the base system itself cannot be updated using this method, but instead, can only be updated by booting from the installation medium, e.g. a CD. Keep in mind that the older the previous version is and the more the package configuration differs from the standard, the more difficult it will be to update it. Under rare circumstances, the old configuration cannot be correctly processed. In this case, configure from scratch. Furthermore, the existing configuration should be backed up before it is updated.

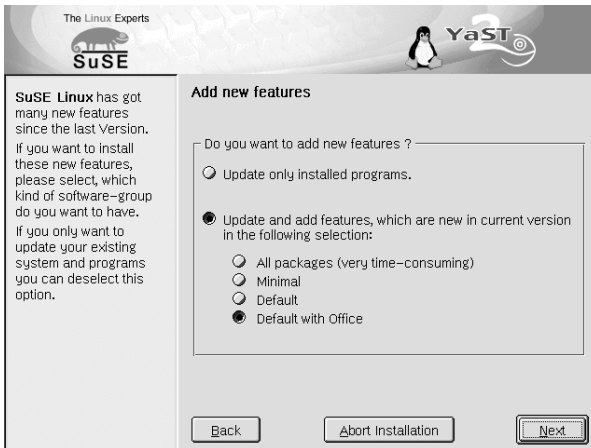


Figure 4.25: YaST2: Updating the System

4.5.7 Boot Mode

The boot mode is normally specified during installation. If you already can boot your SuSE Linux system, you do not need to change anything at this point, unless you have been booting from a floppy and now want to boot from the hard disk.

Otherwise, configuring the boot mode on a running system is only relevant for experts (especially to set kernel parameters after installing a new kernel).

In this dialog, under 'System', define where LILO (**L**inux **L**oader) should be installed. Four options are available to you:

1. 'Write LILO to the boot disk (MBR)'
In the MBR (Master Boot Record) of your hard disk (in `/dev/hda` on IDE systems or `/dev/sda` on pure SCSI systems)
2. 'Create a boot floppy'
3. 'Do not use LILO (a different boot manager is required)'
4. 'Write LILO to a different partition'

If SuSE Linux is the only operating system on your computer, select Option 1, which installs LILO in the MBR of your hard disk. Also choose this option if you want to use LILO as a boot manager for multiple operating systems. First, make sure your operating system can be booted by LILO (usually MS-DOS and Windows 9x/Me). If you are using several operating systems, but are not sure whether they can be booted by LILO, or you want to leave the previous start mechanism unchanged, use the option 'Create a boot floppy'. Thus, you can boot SuSE Linux from the floppy disk.

If you already have a boot manager installed and you want to add SuSE Linux to it, select 'Write LILO to the /boot partition (if you have another boot manager)'. After installing SuSE Linux, reconfigure the existing boot manager and integrate SuSE Linux into the booting process. The items 'Write LILO to a different partition' and 'Kernel boot parameters' are for advanced users. Click on 'Next' to install LILO.

Tip

To install LILO on a boot disk, you do not need to change anything on your previous boot mechanism and can start SuSE Linux from the floppy disk any time. The option 'Create a boot floppy' is therefore the best alternative for the implementation of additional operating systems.



4.5.8 Creating a Boot, Rescue or Module Disk

Using the YaST2 module (under 'System'), create two different types of boot disks, a rescue disk, and two kinds of module disks. Both boot disks enable initial installation if you have problems booting from CD. The disks are actually not intended for booting an already installed system. With a little trick, however, you can still use them to boot an already installed system.

- **Boot disks:** The default boot disk is the one found in your SuSE Linux box. Otherwise, create a boot disk for **i386** and older Cyrix processors.

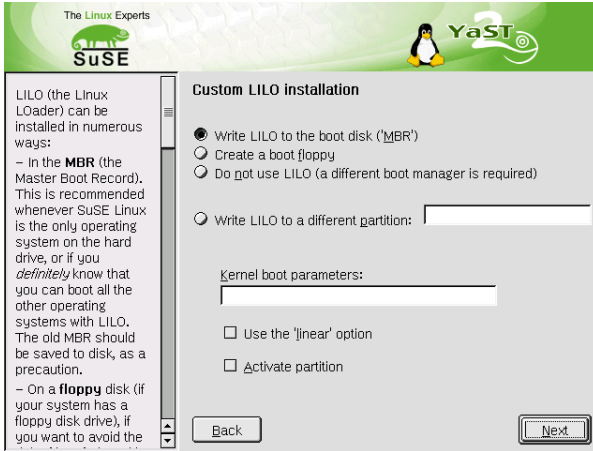


Figure 4.26: YaST2: Configuring the Boot Mode

- **Rescue disk:** The rescue disk can help you regain control access to your system. A “minimal Linux” will be loaded which contains all the helpful tools needed to resolve most problems.
- **Module disks:** If you need additional modules or drivers for your hardware, for example, for installing over the network, create one of these disks:
 - Modules for SCSI/RAID/EIDE and PCMCIA and old CDROM drivers (not for ATAPI)
 - network modules

Select the corresponding item shown on the screen. Insert a (preferably empty or formatted) disk and click ‘Next’. The respective contents will be written to the disk. The above-mentioned boot disks should not be confused with the boot disks used to boot an already installed system. This type of disk will be created, for example, during installation and will start your Linux installed on the hard disk when the floppy is in the drive while your computer is booting.

If all else fails, you can also start an already installed system with the boot disk created above. For this, boot from the floppy disk then, once it asks you to insert the first CD, exit the dialog, to prevent the start of a reinstallation. After making the following language and keyboard entries, you will reach a menu where can choose ‘Start installation / system.’ In the following window, ‘Boot installed system’ will appear.

4.6 Miscellaneous

4.6.1 Hardware Information

YaST2 detects the hardware for the configuration of its components. The technical data it recognizes is displayed in this screen. This is especially useful if you

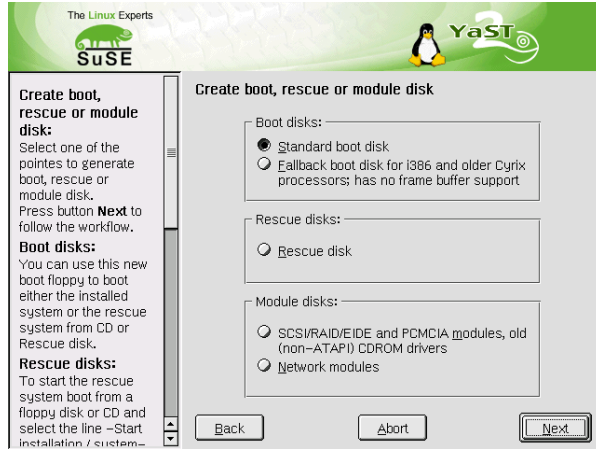


Figure 4.27: YaST2: Creating a Boot/Module Disk

want to post a support request, for instance. You will need hardware information to do this.

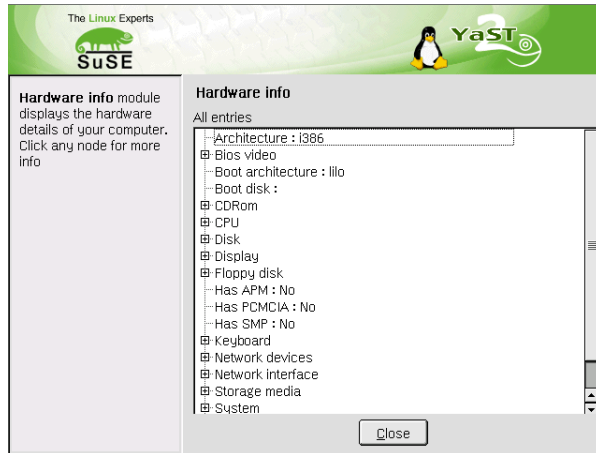


Figure 4.28: YaST2: Displaying Hardware Information

4.6.2 Start Protocol

Start protocol is the screen messages which appear when the system is booting. This protocol is stored in the `/var/log/boot.msg` file. View it easily with this YaST2 module and confirm that all services and functions were started as anticipated.

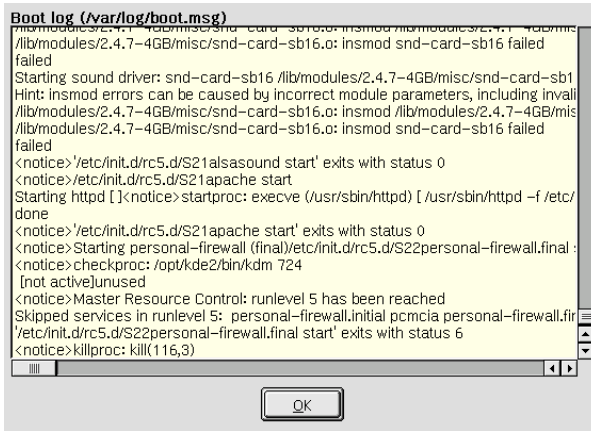


Figure 4.29: YaST2: Displaying the Start Protocol

4.6.3 System Protocol

The system protocol documents the running operation of your computer and is stored in the `/var/log/messages` file. The kernel messages appear here sorted according to date and time.

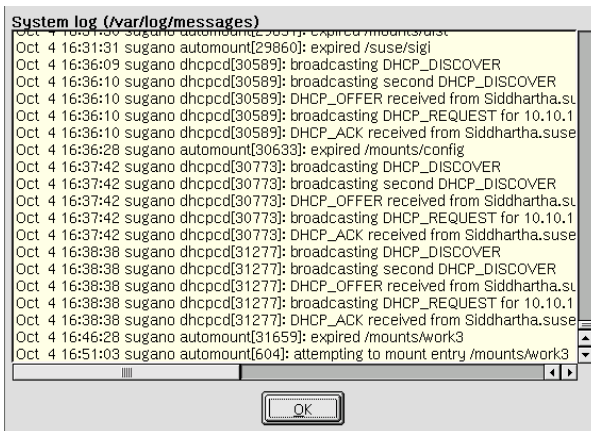


Figure 4.30: YaST2: Displaying the System Protocol

4.6.4 Loading the Vendor's Driver CD

With this module, auto-install the device drivers from a SuSE Linux driver CD.

If you do not need to install your SuSE Linux from scratch, you can load the required drivers from the vendor's CD later with the help of this YaST2 module.

4.6.5 Creating Backups

This option helps you to back up all modified and new files and packages to a file or tape. These are configuration files in most cases.



Figure 4.31: Backup with YaST— Choosing Directories to Exclude

The dialog consists of three parts (see Figure 4.31):

- 1. Choosing the files to back up:
Here, tell YaST which directories should be excluded from the backup. Pre-defined are /tmp, /dev, and /proc. Add mounted CD-ROMs or NFS-mounted file systems to this list. The less you want to be backed up, the faster it will run, since unnecessary comparisons with package lists are omitted. Using (+) and (-), add new directories or remove them. Pressing (F10) leads to the next step.
- 2. Searching:
In this step, YaST searches for files which should be backed up. The number and size of the packages found are updated while searching. After this has been done, there will be a list with all the files that have been found. Here you can still deselect files using (Space).
- 3. Entering commands:
Decide how those files are going to be saved. You can give archive names, options, and more.

This back up mechanism can only work if the dates of the files have not been otherwise changed. Furthermore, this function requires considerable RAM. File names of an ordinary CD take up to 6 MB RAM. Also, you need enough free disk space to save the backup archive. Compressing the archive will lead to a file reduced in size — approximately half of the original. The best way to do backups is to use a tape.

4.7 Important Variables in the rc.config Editor

If you want obtain more functionality for your SuSE Linux Connectivity Server, you can directly control its behavior using the rc.config editor. However, this should only be an exception and is only intended for experts.

After opening the YaST2 rc.config editor, the related variables for the SuSE Linux Connectivity Server will be shown in their own group, all beginning with SLCS.

At present, the following variables can be accessed:

SLCS_SMB_PDC: Should the samba server act as primary domain controller (PDC)?

SLCS_SMB_NAME: The Samba server description

SLCS_PUBLIC_FILESPACE: This is the path to the shared directory (shared volume) . This name must not contain spaces.

SLCS_PUBLIC_FILESPACE_NAME: Description of the shared volume.

SLCS_PRINTER_NAME: Description of our network printer

SLCS_WORKGROUP: Samba workgroup name or domain name

SLCS_ADD_PUBLIC_EXPORTS: Additional directories to export (for nfs, netatalk, smb)

SLCS_NETWORK_DEVICE: Local Intranet network device to be used

SLCS_SQUID_CACHE_DISK: Cache size in MB (on disk) used by the squid HTTP proxy on the hard disk

SLCS_SQUID_CACHE_MEM: Cache size in MB (in RAM) used by the squid HTTP proxy

SLCS_NS_FORWARDERS: Additional name servers

5 Workstation Configuration

In the following chapter we would like to explain how to easily configure Windows and Linux PCs for the use of SuSE Linux Connectivity Servers.

Before going through all the steps of a workstation configuration for SuSE Linux Connectivity Servers data server and Internet functions, please make sure that every PC has a network card on and is connected to the server with the right cable.

There is nothing more frustrating than spending hours searching for the causes of error in applications, just to eventually realize that the network cable was not properly connected ...



Tip

Many network cards are equipped with an LED, which signals that the card is properly connected.

5.1 Access to the File Server

5.1.1 Windows 95/98/ME

Configuration of the Network Card

The first step is complete. The card and cable are connected properly. Now, we just need to make sure that the card is recognized to by the system. If your PC was previously integrated in a network, you can, of course, skip this part and move on to the next step.

Select 'Settings' → 'Control Panel' in the start menu then click on 'System' and select the tab-sheet 'Device Manager' to get an overview of all hardware devices available on your PC.

The item "Network Cards" should now appear in the list, which should indicate the card mounted on your computer, along with its manufacturer and model. Please be aware that items like the "DialUp-Adapter" or "IrDA Infrared Port" in lap-tops have nothing to do with the network card.

In the unexpected event that no such item is listed, you will have to configure the network card first. You will find the necessary information in your network card documentation.

Now close the window "System Properties" by clicking 'OK'.

Installation of Required Components

Once your system detects the network card, make sure that it has all software components needed to allow the Windows system to access the server. In the control panel, select the item 'Network'. It will open a window with three tabs and one overview of the network components installed (see Figure 5.1).

Along with the network card, this list should at least included the installed "Client for Microsoft Networks" and the "TCP/IP-Protocol".

Some PCs have those components. If your PC does not, you will have to install them separately. Select 'Add' and double-click on 'Client'; in the new window, select "Microsoft" as well as "Client for Microsoft Networks" and confirm your selection with 'OK'. For the manufacturer "Microsoft", you will find the "TCP/IP" under 'Protocols'.

Please remember that you will probably need a Windows CD, as well as to restart the system after the installation.

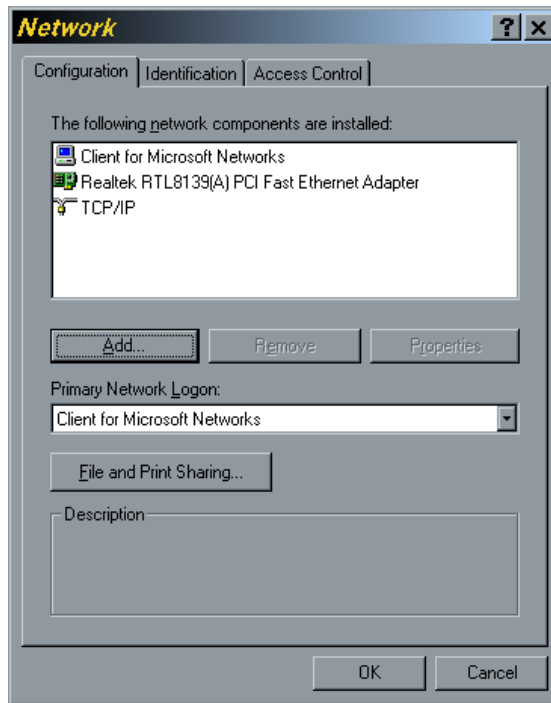


Figure 5.1: Network Configuration in Windows 95/98/ME

Required Settings

After having restarted your computer, all previously installed protocols and services should now be activated. The next step will include a couple of settings which will enable you to have access to the file server.

Double-click on ‘Network’ in the control panel and make first sure that the “Primary Network Registration” is on “Client for Microsoft Networks”. Now select the tab ‘Identification’. Here, you will have to enter some more information. The host name and description are both self-explanatory, except that the former should not contain any more than 15 characters nor spaces.



Note

The most crucial aspect is configuring the work group. The information you type here must correspond to the group that you have already configured during the installation of the SLCS. You could call it, for example, “work-group”. If clients and servers have different work group names, the server cannot be successfully accessed.

If the SuSE Linux Connectivity Server is configured as a domain, you will now have to select the tab ‘Access control’ and change from “share-level access control” to “user-level access control” then indicate the name of the configured domain under “Obtain list of Users and Groups from”. The domain name is derived from the network configuration (see Section 4.3 on page 45, default `slcsnet`)

If you have Windows, enter the corresponding NT domains in the ‘Client for Microsoft networks’ → ‘Properties’ → ‘NT Domains’ menu. The NT domains are taken from the Samba configuration (see Section 4.4.1 on page 48, default `workgroup`).

If your SuSE Linux Connectivity Server was configured as a domain controller (the default setting), all user administration takes place on this server. Every user created is then be recognizable by the Windows clients.

You will not need any further settings, as they will be automatically loaded by the DHCP from the server. As usual, after the configuration of the ID data, you will again have to restart Windows.

If, due to particular system settings, this data cannot be automatically incorporated by Windows, determine whether the IP number of the server has been specified for the Wins/DNS server as well as for the gateway. By default, this is usually 192.168.0.1.

A Preliminary Test

After restarting Windows, you will be asked for a user name and a password. Log in as user with the same characteristics you established in the YaST2 user module. If you have not created any user in the YaST2 Control Center, you can find detailed instructions in Chapter 4.5.1 on page 53.

After booting, click ‘Network Neighborhood’. At most, you will have to wait a few seconds until your SuSE Linux Connectivity Server appears under the names you configured during the installation. Click on the corresponding name and a window will open, containing your private directory (bearing your name), as well as the “Shared Data” area.



Tip

If your login or trial access to the server fails, there might be a network or password problem. A web front-end at <http://password>, along with YaST2 user module on the server, can be used to change your password. This web front-end can even be run from a client host.

If you enter the URL <http://password>, you will be referred to an <https://> address. There are things to keep in mind for this:

1. Note that older clients (Windows 95) often do not provide any support for the https protocol, so in this case, this page will not be accessible.
2. When the connection is first established, the user is informed that this host is not yet known. Normally, the browser will then ask whether this host or key should be accepted. This procedure strictly depends on the browser you are using.

If you are using an older version of Windows 95, which does not support the transfer of coded passwords to Samba, you will have to download an update from <ftp://ftp.microsoft.com/softlib/mslfiles/vrdrupd.exe> and install it to make the encrypted login work.

Linking Drives

Storing data in some deep recess of the network neighborhood would be very complicated indeed. Luckily, you can resort to an easy and convenient solution: linking the network drives of the servers to drive letters.

In 95/98/ME, select the item ‘Extras’ → ‘Link Network Drives’ on Windows Explorer. This will open a window which allows you to link a directory you will have specified under “Path” to a drive letter. For instance, if you link `\\server\public` to `E:`, you will be able to access the contents of the SuSE Linux Connectivity Server public directory as “virtual” drive `E:`.

Please remember that this connection will be canceled at the first shutdown and will not be available at the next start of your Windows system. If you would rather have this drive automatically reconnected the next time the computer is rebooted, select the item “Restore Connection at Next Start”.

5.1.2 Windows 2000

Configuring Windows 2000 for the data service of the SuSE Linux Connectivity Server does not particularly differ from the configuration we performed in the previous chapter for Windows 95/98 or ME. It goes without saying that even Windows 2000 cannot access to the SuSE Linux Connectivity Server without a functioning network card and that even professional Microsoft operating systems depend on certain protocols and settings.



Tip

Please beware, that Windows 2000 requires you to be registered either as “Administrator” or as a user of the “Administrators” group, in order for you to perform such settings of the network configuration.

Configuring the Network Card

Before making sure that all necessary protocols have been installed and all necessary settings completed, we should first check that the network connection card is correctly recognized by the system.

In the start menu, please click on the control panel and select ‘System’. In the “System Properties” window, select the tab ‘Hardware’ and click on ‘Device Manager’, approximately in the middle of the said window. You should find there the details of your network card, together with the name of its manufacturer and model, under “Network Adapter”.

Otherwise, install your network card with the driver, which came with the package, manually, by following the instructions contained in the relevant documentation.

Installing TCP/IP

Just like in Windows 95/98/ME, you will now have to make sure that TCP/IP is installed, so that other services can use it as well.

Double-click on the item “Network and DialUp Connection” in the control panel and select the menu item ‘Properties’ after right-clicking on your network card.

A window will now appear, containing your network card, as well as the “Internet Protocol (TCP/IP)” and the “Client for Microsoft Networks”.

Otherwise, click on the tab “General” and select ‘Protocol’ → ‘Add’. In the dialog ‘Choose Network Protocol’, select ‘Internet Protocol (TCP/IP)’ or ‘Client for Microsoft Networks’ and confirm with ‘OK’.

Make sure that both your network card name and the “Internet Protocol (TCP/IP)”, as well as the “Client for Microsoft Networks” are selected and close the window with ‘OK’.

Verifying Network Identification

Close the device manager and select the tab “Network Identification” in the open “System Properties” window.

Windows 2000 will now show you the name chosen for your PC (usually a combination of letters and numbers), as well as the configured work group or domain. If this information corresponds to the settings on the server (important) and, if

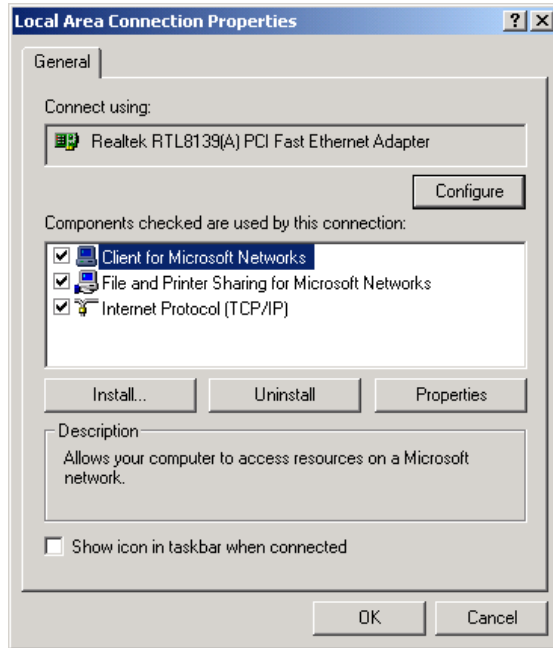


Figure 5.2: Network Configuration in Windows 2000

you are satisfied with the network host name (less important), close this window by clicking OK.

If work groups or domains do not correspond to those of the SuSE Linux Connectivity Server, click the button “Network Authentication”. An assistant will appear, where you will first indicate that this computer belongs to the “Company Network” then click “Next”.

Depending on whether you decided to make use the domain controller functionality of the Samba server, or you prefer to use a work group without a domain, click the option which suits you.

Confirm with “Next”. When you set up a domain, you will be shown a short information page which you will confirm with “Next” again. If, on the contrary, you are configuring a work group, you will only have to add it before completing the setup by clicking on “Finish”. If you decided to use a domain system, you will now also have to indicate your user name and password (as configured on your Linux system) beside the name of the domain.

The Test Can Now Begin

After restarting, the Windows 2000 system should now be properly configured for accessing SuSE Linux Connectivity Server. First, in order to perform a function test, you will need to log in after start-up, with a user name recognized by the server and Windows 2000, along with a valid password. If you have not set

up any users yet, do this now and, if necessary, log in on the Windows 2000 computer with the relevant information.

Tip



The configuration of a user on the SuSE Linux Connectivity Server is described in Chapter 4.5.1 on page 53 in more detail. If you have Windows 2000, you can add a new user in the control panel, under ‘Configure User and Password’.

After logging in, click on “Network Neighborhood” and choose “Neighboring Computers”. In the new window, you should now see your SuSE Linux Connectivity Server underneath the name assigned during installation. Double-click the server name and you will receive a list of all available files, or shares, where you can save your preferred data from now on.

Linking Drives

You will probably want to avoid having to go through the network neighborhood every time you just want to quickly open or save a file on the server. The solution for this is simple: link local drive letters to a network drive.

In order to assign a user-friendly drive letter, such as E:, to a network drive such as \\slcs\SharedVolume, select ‘Workstation’ → ‘Extras’ → ‘Connect Network Drives’ and enter the necessary data.

If you would like the connections to remain active even after the next login, simply select the option “Restore Connection at Next Login”.

5.1.3 With MacOS

In no other operating system is configuring MacOS clients to access your SuSE Linux Connectivity Server so easy. Click on the ‘Apple’ button and select ‘Choice’.

If it is a MacOS 9, you will have to select “AppleShare” before typing the server address. For MacOS X, it will suffice to enter the IP address of your SuSE Linux Connectivity Servers (usually 192.168.0.1). Right after confirming by clicking ‘OK’, MacOS will invite you to enter a user name and password. This login (User name&Password) must correspond to the information saved on the SuSE Linux Connectivity Server.

Tip



If you want to create a new user, you can also make use of the YaST2 Control Center. See Chapter 4.5.1 on page 53.

After a logging in, MacOS will present you with a list of all shares available on

the SuSE Linux Connectivity Server. Choose one and confirm this dialog with 'OK'. The chosen file share will now be available on the desktop.

5.2 Internet Access

One of the most important functions of SuSE Linux Connectivity Servers is undoubtedly the option of conveniently and safely surfing in the Internet using computers in your own local network. In the following sections, we would like to show you how easily you can profit from these possibilities at a workstation.

Proxy Server Setup...

It is generally advisable to use a "proxy server" every time you retrieve Internet data. This program runs on the server and – to put it simply – receives queries from the internal network, retrieves the relevant data from the Internet and transmits them in compressed form back to the internal computer.

A proxy as buffer between your network and the Internet represents a relevant security innovation, as Internet queries (HTTP, FTP) coming from both sides, only arrive so far as the proxy, which will then take over the rest of the transfer.

...with Netscape Communicator

In order to get Netscape to use a proxy server, open the browser and select 'Edit' → 'Settings' on the menu bar. In the window which ensues, select 'Advanced' → 'Proxy'.

Under "Automatic Proxy Configuration File", enter: `http://<Name of your SLCS-Server>/proxy.pac`. This file, stored on the SuSE Linux Connectivity Server, will from now on be read by Netscape each time it starts and contains all relevant proxy settings.

...with Microsoft Internet Explorer

In order to configure Microsoft Internet Explorer for the use of a proxy service, select the item 'Internet Options' in the menu 'Extras'.

Under the tab 'Connections', select 'LAN-Settings' and activate the selection box at "Use Proxy Server". For an address, enter the SuSE Linux Connectivity Server-name (i. e. `server.suse.net`) or the IP address (usually `192.168.0.1`). As "Connection", set 3128.

...with MacOS

MacOS is just as easy to configure for the use of a proxy server. Open the menu "System Settings" and select "Network". Before choosing 'Proxies' in the window in front of you, make first sure that your "Connection" (in the upper part of the window) is "Ethernet" (and not modem). Configure your web and FTP

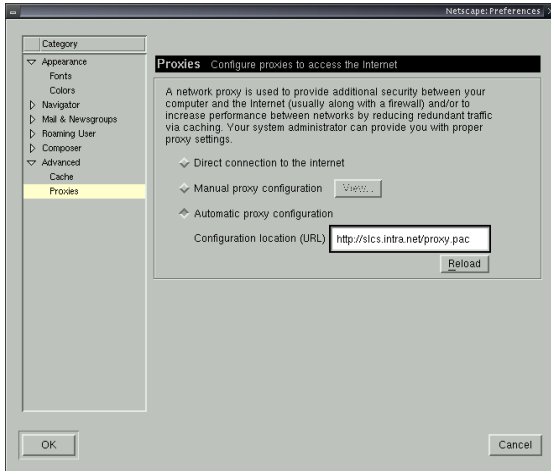


Figure 5.3: Proxy Configuration in Netscape for Linux

proxy as, respectively, 192.168.0.1 as address at the port 3128. We wish you a lot of fun surfing!

Further Settings

If you do not want to use a proxy server but still want to have access to FTP servers on the Internet, then you should activate the option “Use Web-Based FTP” in Internet Explorer of Windows 95/98 or ME, which you will find in ‘Internet Options’ → ‘Advanced’.

Otherwise, if you have Windows 2000, you will need to deactivate ‘Active Folder View for FTP-Sites’, in order to guarantee a perfect FTP data transfer.

Internet Access Controls

If you have configured your SLCS to set up connections to the Internet only when you specify them (“manual connection”) rather than automatically activate connections each time when necessary (“automatic connection”), a special web front-end is available for this purpose. Start your browser (e.g. Netscape Communicator or Microsoft Internet Explorer) and enter the address <http://<NameofyourSLES>/internet.html>. A screen will appear, where connections can be established or terminated whenever you like, just by clicking on the buttons ‘Connect’ and ‘Disconnect’.

5.3 Configuration of SuSE Linux Clients

After having configured several Windows platforms in the previous section, we would now like to show you how to use a SuSE Linux system as client of your

SuSE Linux Connectivity Server.

Basic Configuration

As described above for Windows 95/98/ME/2000, we will first have to make sure that the network card is configured properly. Start the YaST2 Control Center and select the option 'Network Card configuration' from the category 'Network/Basic'. This will open a window where both hardware and software settings for your network card can be defined.

Under "Interface", now select your network card and click on 'Edit'. Make sure that the "Automatic Address Setup (via DHCP)" is active and confirm your changes with "Next", followed by "Finish". Answer yes when the system asks whether or not you would like to save the settings.

Your network card is now configured to automatically adapt itself to the DHCP-server of the SuSE Linux Connectivity Server. Set up the access to file server and user database (NIS) in order to access to the file service of the SuSE Linux Connectivity Servers also from a SuSE Linux system.

File Server Settings

As compared to the previously mentioned Windows solutions, a SuSE Linux system is much easier to configure for accessing your SuSE Linux Connectivity Server. Open the YaST2 Control Center and select the option 'NFS-Client' under 'Network/Advanced'. In the screen which appears, directories from your SuSE Linux Connectivity Servers can be integrated into the file system of the client.

As a standard, the directories /home, where all (personal) home directories of the SuSE Linux Connectivity Server client are stored, and /shared, the public file area accessible to everybody, can both be exported via NFS.

In order to integrate these two new directories, select 'New' and enter the IP address of your SuSE Linux Connectivity Server as "NFS Server Hostname" (usually "192.168.0.1"). Specify /home or /shared as 'Remote Filesystem'.

If /home and /shared are not otherwise required by the client host, we advise you to enter them as mountpoint under 'Mountpoint (local)'. You can, of course, also specify any other directory available on the client. Ignore the 'Options' section if you wish. If necessary, this section may be used by experts to perform some fine-tuning functions. When you are finished with the settings, confirm and save them with 'Finish'.

The file directories of your SuSE Linux Connectivity Server will now be integrated under the mountpoints configured in the Control Center.

NIS Configuration

After installing the NFS, it is advisable to activate an "NIS" as well. NIS is also known as the "Yellow Pages (YP)" and is in charge of ensuring that the logins (i.e. user names and their passwords) of the server are available to every Linux

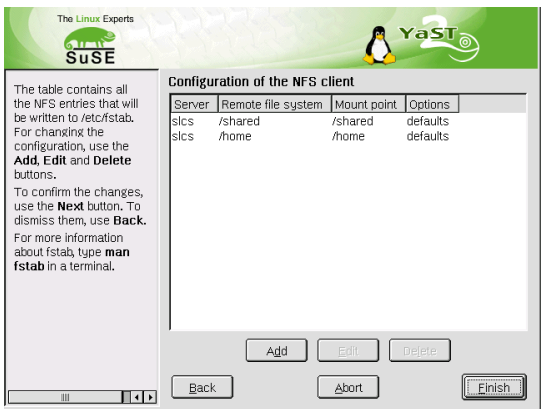


Figure 5.4: YaST2 Module for NFS Configuration

client on your network.

If you would like to have access to private file areas on your server, such as, home directories, you will not be able to do without configuring an NIS, as only authorized users are supposed to have access to this secured area.

To activate an NIS, open the Control Center and select the module 'NIS Client' under the category 'Network/Advanced'. Activate 'Use NIS' and enter the domain configured during the installation in the 'NIS Domain' section (e.g. "suse.net"). The 'IP Address of NIS Server' is usually "192.168.0.1".

Confirm your changes with 'Finish'. Finished. ;)

6 Network Services – Behind the Scenes

This chapter is intended to provide you with more information on network services working for you in the background. Their features are essential for the functioning of your entire network.

6.1 Basic Functions

This section briefly addresses the most basic services required for a functioning network:

Name Resolution The Domain Name Service (DNS) manages names and IP addresses of your local hosts and retrieves the name information from the entire Internet

Configuration of the Network Interface Takes care of assigning IP addresses for your internal clients — by way of DHCP (Dynamic Host Configuration Protocol)

Management and System File Sharing Important user data can be centrally managed and maintained from one main database. The data is exported over NIS (Network Information System)

6.1.1 Domain Name Service

DNS ensures that you never need to memorize an IP address: with the help of DNS, an IP address can be assigned to one or even several names and in turn, a name to an IP address. In Linux, this conversion is usually taken care of by a special software called `bind`. The host responsible for this conversion is called a **name server**.

In doing so, the names comprise a hierarchical system wherein the individual name components are separated by periods. The name hierarchy, however, does not have anything to do with the IP address hierarchy described above.

Let us take a look at a complete name:

```
laurent.suse.de
hostname.domain
```

A complete name — or, as they are referred to in the professional world, a “fully qualified domain name” or **FQDN** — consists of a hostname and a domain segment. The domain segment consists of an arbitrary component — in the above example `suse` — and the **Top Level Domain**, or **TLD**.

Due to historical reasons, TLD assignment is somewhat confusing. In the USA, for example, three-lettered TLDs are used, but in other places, ISO descriptions only consist of two letters. Several TLDs are listed in Table 6.1 to give you an idea.

.com	Commercial — Private companies in the USA.
.edu	Educational — schools, universities and other non-commercial educational institutions in the USA.
.gov	Government — Government institutions and offices in the USA.
.org	Organizational — Non-commercial/non-profit organizations in the USA.
.de	Hosts in Germany.
.at	Hosts in Austria.

Table 6.1: Some Top Level Domains

As you can see, hosts in Germany normally obtain **de**, **at** in Austria, and **ch** in Switzerland.

In the early days of the Internet (before 1990), the names of all the hosts in the Internet were stored in one single file called `/etc/hosts`. However, in light of the rapidly growing number of hosts online, this method was no longer efficient. Therefore, a decentralized and distributed database was designed. The local name server only knows very few of all host names and forwards requests for unknown hosts to other name servers in the Internet.

“Root name servers” can be found at the top of the hierarchy. They manage top level domains. Root name servers are administrated by the Network Information Center, or **NIC** for short. The root name server recognizes each name server responsible for a top level domain. In the case of the German top level domain **de**, the DE-NIC is responsible for domains ending with the TLD **de**. More information on DE-NIC can be obtained at the web site <http://www.denic.de>, and more information on the top level domain NIC can be found at <http://www.internic.net>.

In order for your machine to be able to resolve a name into an IP address, it must be made public by at least one name server with an IP address. The configuration of a name server is easy with YaST2. If you dial in over modem, you might not need to manually configure a name server at all. The protocol needed for the dialup connection is transmitted along with the address of the name server during the dialing process.

Running the Name Server BIND

The name server BIND8, as well as the new version BIND9, is already preconfigured in SuSE Linux, so that you can easily start it up right after you have installed the Linux distribution.

With the configuration files included, your name server will already implicitly recognize all the hosts on your local network. Your name server is able to inform each host on the entire network of your colleague's IP address or fully qualified name.

After you have entered your provider's name server IP address in the YaST2 screen **Hostname % name server configuration** during installation, your name server will also be able to resolve the remaining addresses in the rest of the Internet.

You will know your name server is working when external as well as internal addresses can be resolved using the `host` program.

Further Information

- Documentation on package `bind8`: `file:/usr/share/doc/packages/bind8/html/index.html`.
- A sample configuration can be found at:
`/usr/share/doc/packages/bind8/sample-config`
- The manpage for `named` (**man 8 named**), where the relevant RFCs are named, and in particular, manpage for `named.conf` (**man 5 named.conf**).

6.1.2 DHCP

“Dynamic Host Configuration Protocol” is responsible for creating network settings from a central point on a server, instead of configuring these on all the different workstations. A client configured with DHCP does not have any static addresses of its own, but instead, independently configures them according to the rules set by the DHCP servers.

This enables every client to be identified based on the hardware address of its network card, constantly updated with the same settings, as well as any “interested” host to be “dynamically” assigned addresses out of a certain pool. In this case, the DHCP server will attempt to assign each client the same address for each request (even over a longer time period) — of course, this does not work if there are more addresses on the network than hosts.

System administrators benefit from DHCP in two different ways. For one thing, extensive modifications can even be made to network addresses or to the overall configuration in the DHCP server's central configuration file, without having to configure a large number of clients on an individual basis. Secondly, especially new machines can very easily be integrated into the network, as IP numbers are assigned automatically out of the address pool. The ability to import the appropriate network configuration from a DHCP server is an especially useful feature for laptops, which are often connected to several different networks.

Along with the IP address and the netmask, the client is informed of the host and domain names, as well as the gateway and the name server addresses to be used.

Moreover, numerous other parameters can be configured centrally, such as a time server, from which the current clock time can be queried, or a print server.

Finally, even clients without hard disks (“diskless clients”) can import their operating systems and all their configuration files from the network. However, that is material for a chapter in itself and will therefore only be briefly addressed here.

Further Information

Additional information on DHCP can be found at the web sites the for *Internet Software Consortium* at: <http://www.isc.org/products/DHCP>

Instructions containing advice relating to concrete scenarios can be found in the corresponding man pages:

- General information on the DHCP server daemon: **man dhcpd**
- Information on its configuration:
man dhcpd.conf and **man dhcpd.leases**
- Options to pass to DHCP clients: **man dhcp-options**

6.1.3 NIS

If several UNIX systems have to access common resources in a network, one has to be certain, for example, that user and group names are in sync on all hosts. The network should be transparent for the user. No matter on which host a user is working, he should encounter the same environment. The NIS and NFS services make this possible. NFS is responsible for file system sharing over the network. It is described in more detail in Section 6.2.1 on the facing page.

NIS (Network Information Service) can be characterized as a database service which enables access to information pertaining to important system files network-wide. NIS comes into play by distributing the following files:

/etc/passwd despite this file’s name, it actually only contains data as to login, user ID, group ID, home directly, and the user’s default shell.

/etc/shadow encrypted passwords are located here. In addition, this file contains information on how many days a password is valid.

/etc/group is a list of all network-wide groups which includes the group ID and may also contain optional information as the group’s members.

The advantage of this type of centralized solution: almost all system-wide data only needs to be maintained at one single location in the network. NIS makes any changes public without having to update each host each time a change is made.

Further Information on NIS

Not only can you find information in your own system at `/usr/share/doc/packages/ypbind/`, or in the man pages, “Linux NIS(YP)/NYS/NIS+ HOWTO” can also be found at <http://www.linuxdoc.org/HOWTO/NIS-HOWTO/index.html>.

6.2 File and Print Service

The main job of your server is to manage files and directories as well as print jobs, regardless of what type of operating system the client has. Linux clients are supplied with files and directories over NFS (the “Network File System”). Print jobs can be processed by a printer connected to the network.

Windows clients are connected to your Linux server via Samba, which enables these clients to mount file systems (“shares”) and use the integrated network printer.

6.2.1 NFS – Shared File Systems

As already mentioned in Section 6.1.3 on the preceding page, the purpose of NFS is, along with NIS, to make a network transparent to the user. NFS enables the distribution of file systems over a network. Regardless of which host a user is working on in the network, he will always encounter the same environment. In this manner, by way of NFS, those using your server can have access to your personal home directory without this having to physically exist on any of Linux client machines.

As with NIS, NFS is likewise an asymmetrical service. There are NFS servers and NFS clients. A machine can also be both, in other words, it can simultaneously make file systems available to the network (“export”) as well as mount file systems from other machines (“import”). In the typical scenario, however, servers with a larger hard disk capacity are used for this purpose, and their file systems are usually mounted by clients.

Importing File Systems

Importing file systems from an NFS server is quite simple. The only requirement for this is that the RPC portmapper has to have been started, which is automatically the case following installation. If this condition has been met, remote file systems can be integrated into the file system alongside the local disks by using the command `mount`, provided that the remote file systems are exported by the respective hosts. The syntax is as follows:

```
mount -t nfs <host>:<remote path> <local path>
```

So, for example, use the following command to import user directories on `sun`:

```
earth:~ # mount -t nfs sun:/home /home
```

Exporting File Systems

A host which exports file systems is known as an NFS server. The following network servers must be started on an NFS server:

- RPC portmapper (portmap)
- RPC mount daemon (rpc.mountd)

- RPC NFS daemon (`rpc.nfsd`)

They are started by the `/etc/init.d/portmap` and `/etc/init.d/nfsserver` scripts when booting the system.

After starting these daemons, you must state which file systems to be exported to which hosts. Do this in the `/etc/exports` file.

One line should exist for each directory, indicating which hosts are to access it and how. All subdirectories in these directories will likewise be exported automatically. The permissible hosts are usually indicated by their names (including the domain). The wildcards `'*'` and `'?'` can also be used, which have the same function known in the `bash`. If no host name has been given, each machine may access this directory (with the appropriate permissions).

The permissions to be exported along with the directory are given inside parentheses. The most important options for access permissions are described in the following table.

<code>ro</code>	File system will only be exported read-only (default).
<code>rw</code>	File system will only be exported with write and read permissions.
<code>root_squash</code>	This option causes the user <code>'root'</code> of the specified host not to have the special permissions on this file system which <code>'root'</code> would otherwise have. This is achieved by converting accesses by the root user ID 0 to user ID 65534 (-2). This user ID should be assigned to user <code>'nobody'</code> (default).
<code>no_root_squash</code>	Do not covert root accesses; root permissions are thus retained.
<code>link_relative</code>	Convert absolute, symbolic links (those that begin with <code>'/'</code>) into a relative series of <code>'../'</code> . This option only makes sense when the entire file system of a host is mounted (default).
<code>link_absolute</code>	Symbolic links remain unchanged.
<code>map_identity</code>	The same user IDs are used on the client as on the server (default).
<code>map_daemon</code>	Clients and servers do not have matching user IDs. This option instructs the <code>nfsd</code> to generate a conversion table for the user IDs. The daemon <code>ugidd</code> has to be activated before this entry can be made.

Table 6.2: Access permissions for exported directories

The `exports` file could appear as shown in File Output 6.2.1 on the facing page.

The `/etc/exports` file is read by `mountd` and `nfsd`. Therefore, if this file has been modified, `mountd` and `nfsd` will have to be restarted to apply these changes!

```
#
# /etc/exports
#
/home                sun(rw)   venus(rw)
/usr/X11             sun(ro)   venus(ro)
/usr/lib/texmf       sun(ro)   venus(rw)
/                   earth(ro,root_squash)
/home/ftp            (ro)
# End of exports
```

File 6.2.1: /etc/exports

This is most easily done with the command:

```
earth:~ # rcnfsserver restart
```

6.2.2 Samba

The program package Samba enables you to convert any UNIX machine into a powerful file and print server for DOS, Windows and OS/2 machines. The Samba Project is run by the Samba Team and was originally developed by the Australian ANDREW TRIDGELL.

Samba uses the SMB protocol (Server Message Block) from Microsoft. Due to the initiative of IBM, Microsoft released the protocol so that other software manufacturers could establish connections to a Microsoft domain network. Hosts of other operating systems can communicate with hosts in a Microsoft domain network over SMB. Samba sets the SMB protocol on top of the TCP/IP protocol, meaning that the TCP/IP protocol must be likewise installed on all clients.

Linux can take on the client role as well as the server role — regardless of whether it involves file sharing or print sharing. Samba enables Windows machines to access a Linux file server and to read and store files there. Linux machines can, in turn, also mount file systems distributed by Windows servers (“shares”) and have read and write access to them.

The good thing about this solution: the Linux server, when accessing the network, acts as if it were a Windows machine itself. To put it more precisely: it presents itself to all the integrated Microsoft hosts in the network as a Windows NT 4.2 server. This way, the Microsoft users in the heterogenous network never has the impression that he is on “foreign turf”. Everything appears as it always has, so no costs are incurred by training measures.

Scope of Services

Samba aids Linux computers in providing several services from the Microsoft world. These include:

- File servers
- Print servers

- Primary domain controllers
- Primary WINS servers
- Windows 95/98 authentication

One application does not take care of all these tasks — in Linux, these jobs are done by two “daemons”, or background processes, on the Samba server:

- `smbd` manages resources (file, print and browser services) and is responsible for user authentication as well as SMB data transfer
- `nmbd` is responsible for the name resolution over the NetBIOS and WINS name requests issued by Windows clients

As a Samba client, i.e. a Linux machine, which is to have access to a Windows machine, a Linux client uses the following programs:

- `smbclient` enables Windows file system access
- `smbtar` lets you save SMB shares to Unix tape drives
- `nmblookup` name resolution for NetBIOS names
- `smbpasswd` SMB user password management
- `smbstatus` information on open SMB connections

All these applications are *Samba Suite* components and more or less work for you in the background.

Further Information

Meanwhile, there are books and web pages which handle the topic of Samba — but there is also quite a bit of useful and relevant information on Samba available. Take a look at `/usr/share/doc/packages/samba/`. You will be greeted with a wealth of information here. In addition, the complete version of the book “Using Samba” by Robert Eckstein, David Collier-Brown, and Peter Kelly is located under `/usr/share/doc/packages/samba/htmldocs/using_samba/`.

Furthermore, the Samba project web sites have something for you:

`http://de.samba.org/samba/samba.html` is the official mirror of the Samba site. In the subdirectory `http://de.samba.org/samba/docs/`, you will find a summary of the most important (and current) information sources (including man pages) on this topic.

6.3 Security

6.3.1 Firewall

Your server is protected from attacks coming from the Internet by an easy-to-use packet filter. The personal firewall works almost maintenance-free and ready for

implementation following a single configuration step. When it is active, it allows access to the Internet from the inside out but blocks connections from the outside in.

Since the SuSE Linux Small Business Server was conceived as a pure file and print server for private networks and does not offer any Internet services (FTP, HTTP etc.), it can be easily but very effectively secured with this solution.

When the personal firewall is active, all data packets belonging to any one of these three categories will be refused:

- UDP packets
- Attempted external TCP requests
- ICMP Redirect Subtypes (ICMP Redirects can be used to trigger your machine to change its routing table)

Personal Firewall is exclusively configured using a single variable stored in the file `/etc/rc.config.d/security.rc.config`.

The variable which needs to be configured is `REJECT_ALL_INCOMING_CONNECTIONS`. As soon as a reasonable configuration option has been chosen for this, the firewall will automatically start after configuration is completed and the network rebooted.

The following configurations are possible:

no	If "no" is specified — or if this field is left blank, the Personal Firewall will not be active. All incoming connections will be accepted. No filtering takes places.
yes	Personal Firewall affects all interfaces other than "lo", the loopback interface, "localhost". Thus, connections originating from inside the network will be blocked. The individual packets which are accepted are those directed to "localhost".
iface	Here, interfaces on which connections should be blocked are explicitly given here (and divided by blank spaces).
masq	Packets which make it to a machine but is not permissible for one of its interfaces will be masked before being forwarded. Here, the name of the interface is specified over which masked packets are to gain access to the outside and where incoming connections are to be rejected. (Interface name and "masq" must be separated from one another by a blank space.)

Table 6.3: Configuration of the Personal Firewall

When configuring your Internet access via Modem or ISDN, the YaST2 screens 'Connection parameters' and 'ISDN connection parameters' will ask you whether the firewall should be activated. Selecting this option is the same as entering "masq Interface Name" in the `/etc/rc.config.d/security.rc`.

config file. By way of masquerading, all network packets coming from internal clients destined for the Internet will not be tagged with their originating network address, but appear as if they originated the network interface on your server recognized in the Internet. On one hand, your internal network received additional protection in that the individual clients are only locally known, and on the other, this saves space for Internet addresses. If you choose not to activate the firewall, network traffic on your Internet connection will not be filtered at all (synonymous with a "no" in the configuration file).

6.4 Proxy Server: Squid

The following chapter describes how caching web sites assisted by a proxy server works and what the advantages of using Squid are.

Squid is the most popular proxy cache for Linux/UNIX platforms.

What is a Proxy Cache?

Squid acts as a proxy cache. It behaves like an agent which receives requests from clients (in this case web browsers) and passes them to the specified server provider. When the requested objects arrive at the agent, it stores a copy in a disk cache.

Benefits arise when different clients request the same objects: these will be served directly from the disk cache, much faster than obtaining them from the Internet and, at the same time, saving overall bandwidth from the system.

Tip

Squid covers a wide range of features including intercommunicating hierarchies of proxy servers to divide the load, defining strict access control lists to all clients willing to access the proxy and, with the help of other applications, allowing or denying access to specific web pages. It also can obtain statistics about the most visited web sites, user usage of the Internet, and many others.



Squid is not a generic proxy. It proxies normally only between HTTP connections. It does also support the protocols FTP, Gopher, SSL and WAIS, but it does not support other Internet protocols such as Real Audio, news or videoconferencing. Because Squid only supports the UDP protocol to provide communication between different caches, many other multimedia programs will not be supported.

Squid and the SuSE Linux Small Business Server

Squid will already run when you boot your system for the first time, without having to do anything on your part. The basic functions it takes care of for you are:

Web Site Caching All clients on the internal network can benefit from caching the requested web sites.

Restricted access inside the network The configuration of Squid is structured in such a way that only local clients have access to its services.

Cache Management If you have an Apache web server set up, the Cache Manager application allows you to query current statistics on your server this at any time, regarding the amount of memory required for Squid's caching functions.

You can supplement this basic functionality with some other useful features by modifying the configuration file `squid.conf` accordingly. Reasonable extension are, for example:

Refined Internet access rules by way of ACLs, or Access Control Lists, limit Internet access for particular user groups only to certain times of day, for example. Here, an example:

```
acl mysurfer srcdomain .my-domain.com
acl teachers src 192.168.1.0/255.255.255.0
acl students src 192.168.7.0-192.168.9.0/255.255.255.0
acl afternoons time MTWHF 12:00-15:00

...

http_access allow localhost
http_access allow teachers
http_access allow students afternoon
http_access deny all
```

File 6.4.1: Excerpt from a `squid.conf` with Access Restrictions

This configuration allows the user group 'teachers' unlimited Internet access at any time of the day, the user group 'students' only during the afternoon, and all other users, no access at all.

Internet access only for authenticated users If you only want to allow access for authorized users, integrate an authentication program such as `pam_auth`, which asks every user for his login and password:

```
authenticate_program /usr/sbin/pam_auth
...

acl password proxy_auth REQUIRED

http_access allow password
http_access deny all
```

File 6.4.2: Proxy Authentication in `squid.conf`

Additionally, another ACL will have to be set up so that only clients with a valid login can surf. Alternatively, `REQUIRED` can be substituted by a list of permitted user names.

ID Request Provided you have the right software, you can also configure Squid so that it asks every user who wants to surf for his ID. In Linux, the program `pident` is used for this purpose. The related software is available in the Internet as a free download for Windows clients.

Blocking undesired URLs Undesired URLs can be blocked for certain users by way of properly set up ACLs and with a separate program, such as SquidGuard, in conjunction with proxy authentication. *Contents*, or certain *script languages* embedded in HTML (Java Script, VBscript), however, cannot be filtered, censored or blocked by Squid or SquidGuard in any way.

Fine-tuning With a little experience, you can further improve your proxy's performance by synchronizing the desired cache size with the memory capacity of your system. Using multiple caches also conserves your system's memory, since objects can then be swapped with other caches.

More Information on Squid

Visit the home page of Squid: <http://www.squid-cache.org/>. Here, you will find well-versed information on the topics of "Multiple Cache Usage", "Tuning the Cache", "Setting Up ACLs" and "Setup of the Configuration file" (<http://squid.visolve.com/squid24sl/contents.htm> is especially recommended for this, and includes extensive explanations of all the different configuration options), and much more.

Information on using a cache manager can be found at: <http://www.squid-cache.org/Doc/FAQ/FAQ-9.html>

The Calamaris is also relevant in this context. Calamaris is a Perl script which generates cache reports in HTML or in ASCII format: <http://Calamaris.Cord.de>

If you need information on SquidGuard, you can find all kinds of useful tidbits at the project home page:

- General information at: <http://www.squidguard.org>
- Sample configuration and explanations at:

<http://www.squidguard.org/config/>

Furthermore, mailing lists for Squid can be found at: squid-users@squid-cache.org.

The mailing list archive is located at:

<http://www.squid-cache.org/mail-archive/squid-users/>

6.5 Intranet Server

In conclusion, a few words about the Intranet server, which will run on your system — Apache. This will provide essential services for network—internal presentations of some services and also, internal documents.

6.5.1 Apache

A — if not *the* prestigious project of the open-source scene is the web server Apache. About 60% of all web servers around the world run on this *free* software, quite obviously underscoring the fact that open source products can indeed operate smoothly in the professional arena.

Apache was originally conceived as a type of makeshift solution for enhancing the NSCD 1.3 web server, featuring important innovations and bug fixes. Hence the name (“a patchy server”), which has more to do with its “patchwork structure” than with a North American Indian tribe.

Apache as Intranet Server

The following section will basically provide you with a compact description of how to configure the Apache running on your system to work for you as an Intranet server. Please understand that we cannot offer you a complete explanation of all the configuration options and supplementary modules. For this information, refer to the information sources on page 95.

The Apache configuration file is nearly 1500 lines long and not necessarily intended to enthrall the newcomer with intuitive user-friendliness. However, the default configuration provided with the application package is more than sufficient for most purposes. The selected default settings are designed so that the web server can run smoothly on most systems straight “out-of-the-box”.

In order to derive a powerful Intranet server from a newly installed Apache, the following steps are required:

1. Which contents should be presented?

The directory, from which Apache typically awaits the contents to be presented, is `/usr/local/httpd/htdocs`. Store all files here, which are to ultimately be displayed by Apache — and with read permissions for other users:

```
-rw-r-r- 1 me my_group 0 Mar 2010 14:26 my_file
```

Directories are created with read, write and execute permissions for you, and read-execute permissions for your group and others. SuSE Apache expects your own personal home page to be called `index.html`. If you do not have your home page under this name, Apache will access its own default home page.

2. Now, the most basic entries can be made (as ‘root’) in the `/etc/httpd/httpd.conf` file. These include:

ServerRoot /usr/local/httpd the file tree, in which all Apache-specific files are located, are stored here.

ServerAdmin the administrator's e-mail address — if a page cannot be displayed on the new server, this address will be relayed along with other data

ServerName fully qualified name of the server

DocumentRoot the “DocumentRoot” variable is /usr/local/httpd/htdocs, as specified in Step 1

CustomLog /var/log/httpd/access_logcommon. By default, access attempts are logged on your server here.

ErrorLog /var/log/httpd/error_log, error messages are archived here

3. If individual directories or even the contents of the server are to be protected from unauthorized “entry”, there are several simple security mechanisms which can be implemented for this:
 - Protection of specific directories via .htaccess. Every directory you want to have secured receives its own .htaccess file, containing the following relevant lines:

```
order deny,allow
deny from all
allow from slcs.com
```

The directory (and all subdirectories, where this file is located, is now secured against external access attempts and only allows the sending of data to hosts in the internal domain slcs.com. You can proceed in the same manner in order to block certain hosts from the slcs.com domain, even if the Internet connection is fully insulated by a firewall. For the newly created .htaccess file to be read at the start-up of Apache, the following option must be enabled in the /etc/httpd/httpd.conf file:

```
#
# This controls which options the .htaccess files in directories
# can override. Can also be "All", or any combination of
# "Options", "FileInfo", "AuthConfig", and "Limit"
#
AllowOverride All
```

If the option above is not enabled, all .htaccess entries will be ignored.

- The same effect as in .htaccess is also achieved by making the parallel entries in the main httpd.conf file. Either individual <Directory/> entries are specified or permission restrictions are set, according to category:

This method has the same effect as if you were to secure the highest directory level of the web server, the **DocumentRoot**, including all its sub-directories by way of .htaccess.

```
#  
# Controls who can get stuff from this server.  
#  
    Order deny,allow  
    Deny from all  
    Allow from slcs.com
```

This brief overview serves as an introduction to the basics of the security of directories, rather than to serve as a set of detailed instructions. In addition, a large selection of other complex security mechanisms exists, which can be implemented by an administrator to shake off unwanted visitors. In the same vein, certain regions on a server can likewise be restricted to only authorized users with the right password.

Further Information on Apache

Find out more about Apache at the project web site: <http://httpd.apache.org>. Quite extensive information as to the current state of development, FAQs, tutorials, and an excellent explanation of its configuration is available here .

If you are interested in additional modules for your web server, <http://modules.apache.org> is a good place to start.

For extremely in-depth questions and higher standards, several books have been published, most notably the O'Reilly publication "Apache: The Definitive Guide" by Ben and Peter Laurie.

A weekly newsletter is published in ApacheWeek (<http://www.apacheweek.org>), providing information on the most current developments in the project.

7 ALICE

The automatic installation and configuration of Linux systems enables the setup of a unified server landscape. The automatic method is even preferable to manual installation and configuration for clients — beyond a certain number of units. This standardization affects system and software versions, file system structures, and configuration files.

Automation guarantees that a previously determined and successful installation method can be reapplied on a machine at any time, even without professional knowledge. Thus, expanding the network landscape is also made easy. Standardization simplifies the administration as well: given the same configuration structures, configuration files do not have to be tracked down again and again on each server. The software behaves in a predictable manner, due to its equal version status. Bug fixes prevail over the entire network landscape.

All these aspects contribute to improved production quality.

7.1 What is ALICE?

ALICE ("Automatic Linux Installation and Configuration Environment") is an application package consisting of various modules which are responsible for installing and configuring workstations and servers. This is carried out in two steps. In the first step, configuration files and the boot medium are generated. The second step proceeds with the installation and configuration of the machines, in accordance with the configuration files.

With ALICE, you can install individual machines, such as servers, as well as numerous related machines (cluster nodes, server farms).

To simplify the configuration process, a different configuration is not created for each individual machine. Instead, the machines can be classed. When using this classing system, only the unique characteristics of each machine have to be specified (such as the network IP address).

Some basic understanding of system administration in Linux/Unix is required to be able to use ALICE.

7.2 How ALICE is Installed

ALICE is needed to prepare the configuration on a workstation and the target system to be installed. Of course, ALICE is installed automatically on the target system and package `alice`, series `n`, on the workstation.

The configuration is stored under the root directory in `$ALICE_HOME`, e.g. `/home/myuser/projects/alice`. As previously mentioned, the machines are categorized into classes while retaining different distinguishing characteristics. For this reason, three directories are located in `$ALICE_HOME`:

1. The `classes` directory. All the classes are stored here.
2. The `info` directory. Special settings for each machine are stored here.
3. The `templates` directory. The default settings are found here.

A sample configuration can be found under `/usr/lib/alice/samples`.

In order to maintain a better overview of the configuration, settings for the classes as well as for the individual machines are divided up into three sections, including `sys`, `network` etc. All the different settings are then entered into related `tcf` files, structured as shown in the following

```
<TAG1> .... </TAG1>
<TAG2>
....
</TAG2>
```

File 7.2.1: Basic Structure of an ALICE Configuration File

where anything between `<TAG> .. </TAG>` is seen as a value, including all special characters. The name of the `tcf` file is comprised of: `<classname/hostname>.<section>.tcf`

7.3 Creating a Simple Configuration

The speed and simplicity of installing a machine with ALICE is well illustrated by the following simple example.

These steps are to be carried out by `'root'`, because only `'root'` has the necessary permissions to do so. This example only shows one method of installing with ALICE and places the least amount of demand on the infrastructure. ALICE also enables you to install without using any floppy disks at all, but the disk-less method does burden the infrastructure somewhat and it requires NFS, DHCP, and TFTP servers as well as PXE-compatible network cards, or a PXE or NETBOOT boot floppy.

In the typical scenario, one server functions as an installation and configuration server. The various machines are then booted using a boot floppy, and then installed.

The procedure is as follows:

First, the environment variable `$ALICE_HOME` must be set to `/usr/lib/alice/samples` (in the bash with `export ALICE_HOME=/usr/lib/alice/samples`). If the machine has a disk smaller than 6 GB, the file `/usr/lib/alice/sample/`

```
<SYS_PART_hda>
/      6000 num=1 fsys=reiser
SWAP   256 num=2
</SYS_PART_hda>
```

info/simple.sample.de.sys.tcf must be edited, substituting 6000 for the disk size minus 256 MB SWAP:

If the CD-ROM is not /dev/hdc (Secondary master on IDE bus), the following tag must be changed:

```
<SYS_CDROM_DEVICE>/dev/hdc</SYS_CDROM_DEVICE>
```

If the network is to also be configured, both of the following tags must be modified:

- Specify module name of the network driver:

```
<SYS_INSMOD_MODULES>tulip</SYS_INSMOD_MODULES>
```

Other network card modules include 8139two, eepro100, ne2000.

- Basic IP configuration

This tag is defined in the "network section" — in our example, in the file /usr/lib/alice/sample/info/simple.sample.de.network.tcf.

There is one line per interface which looks like

```
<Interface> <IP address> <netmask> <broadcast address>
```

An example:

```
eth0    192.168.1.200    255.255.255.0 192.168.1.255
```

After these adjustments have been made, the respective boot floppy is created by first mounting the CD-ROM/DVD to /media/cdrom as 'root' then the floppy created with

```
/usr/lib/alice/util/make_inst_disk --url file:/media/cdrom simple.sample.de
```


Afterwards, another configuration disk must be created. This is done with

```
/usr/lib/alice/util/make_config_disk /dev/fd0
```

After generating both floppies, installation of the new machine may begin:

- Insert the boot floppy into the drive.
- Insert the CD-ROM/DVD.
- Make sure you are not booting from CD-ROM/DVD (note the boot sequence in the BIOS settings).
- Wait until YaST asks you whether partitioning should be performed. If you confirm this step with yes, the hard disk will be reformatted.

```
<NET_IP_CONFIG>
eth0      10.70.132.33 255.255.0.0  10.70.255.255
</NET_IP_CONFIG>
```



Caution

Normally, repartitioning will cause all previously saved data to be lost.

- YaST now carries out the installation. The boot floppy should now be removed and the configuration disk inserted in its place.
- Following that, wait until all packages have been installed and you have been asked for the root password. The root password can also be automatically set by ALICE by inserting the following tags in the `simple.sample.de.sys.tcf` file:

```
<SYS_SET_ROOT_PWD>yes</SYS_SET_ROOT_PWD>
<SYS_ROOT_START_PWD>laM8LehhunciE</SYS_ROOT_START_PWD>
```

(The password will then be "blank")

- YaST may prompt you to insert additional CDs for installing the remaining packages.
- Finally, the system will reboot and ALICE will finish the configuration.

Voilà — the new system is installed and configured.

ALICE provides many other options as well, above and beyond those described in the example presented above.

7.4 Further Information

Here, you will find further information about ALICE:

<http://www.suse.de/~fabian>

<http://list2.suse.com/alice>

8 Security and Confidentiality

8.1 Basic Considerations

One of the main characteristics of a Linux or UNIX system is its ability to handle several users at the same time (multiuser) and to allow these users to perform several tasks (multitasking) on the same computer simultaneously. Moreover, the operating system is network transparent. The users often do not know whether the data or applications they are using is provided locally from their machine or made available over the network.

With the multiuser capability the respective data of different users must be stored separately. Security and privacy need to be guaranteed. “Data security” was already an important issue, even before computers could be linked through networks. Just like today, the most important concern was the ability to keep data available in spite of a lost or otherwise damaged data medium, a hard disk in most cases.

This chapter is primarily focused on confidentiality issues and on ways to protect the privacy of users, but it cannot be stressed enough that a comprehensive security concept should always include procedures to have a regularly updated, workable, and tested backup in place. Without this, you could have a very hard time getting your data back — not only in the case of some hardware defect, but also if the suspicion arises that someone has gained unauthorized access and tampered with files.

8.2 Local Security and Network Security

There are several ways of accessing data:

- Personal communication with people who have the desired information or access to the data on a computer
- directly from the console of a computer (physical access)
- over a serial line
- using a network link

In all these cases, a user should be authenticated before accessing the resources or data in question. A web server might be less restrictive in this respect, but you still would not want it to disclose all your personal data to any surfer out there. On a SuSE system, a few tweaks are sufficient to make it boot right into your

desktop without even asking for a password, but, in most cases, that would not be such a good idea, as anybody could change data or run programs.

In the list above, the first case is the one where the highest amount of human interaction is involved, such as when you are contacting a bank employee and are required to prove that you are the person owning that bank account. Then you will be asked to provide a signature, a PIN, or a password to prove that you are the person you claim to be. In some cases, it might be possible to elicit some intelligence from an informed person just by mentioning known bits and pieces here and there to win the confidence of that person by using clever rhetoric. The victim could be led to gradually reveal more information, maybe without even becoming aware of it.

Some people are rather unmindful of what they say or act unconsciously in the way they give answers, so that even a question which they believe was left unanswered might provide enough information to proceed with an even more precise question. Piece after piece gets added to the puzzle until the picture is nearly complete (“No, Mr. Smith is on vacation right now, it’s at least three weeks before he’ll be back in. He’s not my boss anyway, you know he’s up there in the fourth floor while I’m here in the third!”). Among hackers, this is called “social engineering”. You can only guard against this by educating people and by dealing with language and information in a conscious way. Before breaking into computer systems, attackers often try to target receptionists, service people working with the company, or even family members and, in many cases, such an attack based on social engineering will only be discovered at a much later time.

A person wanting to obtain unauthorized access to your data could also use the traditional way and try to get at your hardware directly. Therefore, the machine should be protected against any tampering so that no one can remove, replace, or cripple its components. This also applies to backups and even any network cable or the power cord. Likewise, it might be necessary to secure the boot procedure, as there are some well-known key combinations which invoke special reactions during booting. Protect yourself against this by setting passwords for the BIOS and the bootloader.

Serial terminals connected to serial ports are still used in many places, but are rarely installed with new systems anymore. With regard to data access, serial terminals are a special case. Unlike network interfaces, they do not rely on a network protocol to communicate with the host. A simple cable, or maybe an infrared port, is used to send plain characters back and forth between the devices. The cable itself is the weakest point of such a system: with an older printer connected to it, it is really easy to record anything that runs over the wires. What can be achieved with a printer can also be accomplished in other ways, depending on the effort that goes into the attack.

Networks make it easier for us to access data remotely, but they do this with the help of network protocols which are often rather complex. This might seem paradoxical at first, but is really indispensable if you wish to remotely control a computer or to retrieve data from it, no matter where you are. It is necessary to have abstract, modular designs with layers that are more or less separate from each other. We rely on such modular designs in many daily computing situations. Modularity means that your text processor, for example, does not need

to know about the kind of hard disk you use or your e-mail program should not be concerned with whether you have a modem or an ethernet card. Components of your operating system, Linux in this case, provide the necessary functions and make these available to the system through a predefined interface. With this modularity, a text processor or a mail user agent (MUA) can function on a variety of hardware platforms and you can run them from some place in the world with the necessary equipment.

Regarding the data, there is no difference between opening a file from a command line or looking at it with a web browser. The file could also be read via a network (using a telnet program or with a secure shell client — which is actually a much better option as ssh encrypts all network traffic). To do so, the host and the network need to be connected and the user needs to log in and authenticate. The possible actions are still restricted, however, by the file permissions.

Reading a file locally on a host requires other access rules than opening a network connection with a server on a different host. There is a distinction between local security and network security. The line is drawn where data has to be put into packets to be sent somewhere else.

8.2.1 Local Security

Local security starts with the physical environment in the location where the computer is running. Assume that your machine is set up in a place where security is in line with your expectations and needs.

The main goal of “local security” is to keep users separate from each other, so that no user can assume the permissions or the identity of another one. This is a general rule to be observed, but it is especially true for the user `'root'` who holds the supreme power on the system. User `'root'` can take on the identity of any other local user without being prompted for the password and read any locally stored file.

For an attacker who has obtained access to local resources from the command line, there is certainly no shortage of things that could be done to compromise the system.

Passwords

On a Linux system, passwords are, of course, *not* stored as plain text and the text string entered is not simply matched with the saved pattern. If this were the case, all accounts on your system would be compromised as soon as someone got access to the corresponding file. Instead, the stored password is encrypted and, each time it is entered, is encrypted again and the two encrypted strings are compared. Naturally, this will only work if the encrypted password cannot be reverse-computed into the original text string. This is actually achieved by a special kind of algorithm, also called “trapdoor algorithm,” because it only works in one direction. An attacker who has obtained the encrypted string will not be able to get your password by simply applying the same algorithm again. Instead, it would be necessary to test all the possible character combinations until

a combination is found which looks like your password when encrypted. As you can imagine, with passwords that are eight characters long, there are quite a number of possible combinations to calculate.

In the seventies, it was argued that this method would be more secure than others due to the relative slowness of the algorithm used, which took a few seconds to encrypt just one password. In the meantime, however, PCs have become powerful enough to do several hundred thousand or even millions of encryptions per second. Because of this, encrypted passwords should not be visible to regular users (`/etc/shadow` cannot be read by normal users). It is even more important that passwords are not easy to guess, in case the password file becomes visible due to some error. Consequently, it is not really useful to “translate” a password like “tantalise” into “t@nt@1ls3”.

Replacing some letters of a word with similar looking numbers is not safe enough. Password cracking programs which use dictionaries to guess words also play with substitutions like that. A better way is to make up a word with no common meaning, something which only makes sense to you personally, like the first letters of the words of a sentence or the title of a book, such as “The Name of the Rose” by Umberto Eco. This would give the following safe password: “TNotRbUE9”. By contrast, passwords like “beerbuddy” or “jasmine76” are easily guessed even by someone who has only some casual knowledge about you.

The boot procedure

Configure your system so it cannot be booted from a floppy or from CD, either by removing the drives entirely or by setting a BIOS password and configuring the BIOS to allow booting from a hard disk only.

Normally, a Linux system will be started by a boot loader, allowing you to pass additional options to the booted kernel. This is crucial to your system’s security. Not only does the kernel itself run with root permissions, but it is also the first authority to grant root permissions at system start-up. Prevent others from using such parameters during boot by using the options “restricted” and “password=your_own_password” in `/etc/lilo.conf`. Execute the command **lilo** after making any changes to `/etc/lilo.conf` and look for any unusual output the command might produce. If you forget this password, you will have to know the BIOS password and boot from CD to read the entry in `/etc/lilo.conf` from a rescue system.

File Permissions

As a general rule, always work with the most restrictive privileges possible for a given task. For example, it is definitely not necessary to be ‘root’ to read or write e-mail. If the mail program you use has a bug, this bug could be exploited for an attack which will act with exactly the permissions of the program when it was started. By following the above rule, minimize the possible damage.

The permissions of the more than 200,000 files included in a SuSE distribution are carefully chosen. A system administrator who installs additional software

or other files should take great care when doing so, especially when setting the permission bits. Experienced and security-conscious system administrators always use the `-l` option with the command `ls` to get an extensive file list, which allows them to detect any wrong file permissions immediately. An incorrect file attribute does not only mean that files could be changed or deleted. These modified files could be executed by `'root'` or, in the case of configuration files, that programs could use such files with the permissions of `'root'`. This significantly increases the possibilities of an attacker. Attacks like this are called cuckoo eggs, because the program (the egg) is executed (hatched) by a different user (bird), just like a cuckoo would trick other birds into hatching its eggs.

A SuSE Linux system includes the files `permissions.easy`, `permissions.secure`, and `permissions.paranoid`, all in the directory `/etc`. The purpose of these files is to define special permissions, such as world-writable directories or, for files, the `setuser` ID bits, which means the corresponding program will not run with the permissions of the user that has launched it, but with the permissions of the file owner, `'root'` in most cases. An administrator may use the file `/etc/permissions.local` to add his own settings. The variable `PERMISSION_SECURITY`, set in `/etc/rc.config`, defines which of the above files is used by SuSE's configuration programs to set permissions accordingly. As a more convenient way to select the files, use the submenu `'Security'` in YaST1 or YaST2. To learn more about the topic, read the comments in `/etc/permissions` or consult the manual page of `chmod` (`man chmod`).

File race conditions

Assume that a program wants to create a file in a directory which is world-writable (such as `/tmp`). First, the program checks whether the file already exists and, if that is not the case, creates it. However, between checking and file creation, there is a short moment which can be used by an attacker to create a symbolic link, a pointer to another file. The program may then be tricked into following the symbolic link, overwriting the target file with its own permissions. This is called a race because the interval during which the attacker can create a "symlink" is very short. The race is only possible if the checking and file creation procedure is not atomic (indivisible). If the race is allowed to take place at all, there is a chance that it may be won by the attacker. It is all a matter of probability.

Buffer overflows and format string bugs

Special care must be taken whenever a program is supposed to process data which can or could be changed by a user, but this is more of an issue for the programmer of an application than for regular users. The programmer has to make sure that his application will interpret data in the correct way, without writing them into memory areas that are too small to hold them. Also, the program should hand over data in a consistent manner, using the interfaces defined for that purpose.

A “buffer overflow” can happen if the actual size of a memory buffer is not taken into account when writing to that buffer. There are cases where this data (as generated by the user) uses up some more space than what is available in the buffer. As a result, data is written beyond the end of that buffer area, which, under certain circumstances, makes it possible that a program will execute program sequences influenced by the user (and not by the programmer), rather than just processing user data. A bug of this kind may have serious consequences, in particular if the program is being executed with special privileges).

“Format string bugs” work in a slightly different way, but again it is the user input which could lead the program astray. In most cases, these programming errors are exploited with programs executed with special permissions — `setuid` and `setgid` programs — which also means that you can protect your data and your system from such bugs by removing the corresponding execution privileges from programs. Again, the best way is to apply a policy of using the lowest possible privileges.

Given that buffer overflows and format string bugs are bugs related to the handling of user data, they are not only exploitable if access has been given to a local account. Many of the bugs that have been reported can also be exploited over a network link. Accordingly, buffer overflows and format string bugs should be classified as being relevant for both local and network security.

Viruses

Contrary to what some people will tell you, there *are* viruses that run on Linux. However, the viruses that are known were released by their authors as “proof of concept”, meaning that they were written to prove that the technique works as intended. On the other hand, none of these viruses have been spotted “in the wild” so far.

Viruses would not be able to survive and spread without a host on which they can live. In our case, the host would be a program or an important storage area of the system, such as the master boot record, which needs to be writable for the program code of the virus. Owing to its multiuser capability, Linux can restrict write access to certain files, which is the case especially with system files. Therefore, if you did your normal work with `'root'` permissions, you would increase the chance of the system being infected by a virus. By contrast, if you follow the principle of using the lowest possible privileges as mentioned above, chances of getting a virus are slim. Apart from that, you should never rush into executing a program from some Internet site that you do not really know. SuSE's RPM packages carry a cryptographic signature as a digital label that the necessary care was taken to build them. Viruses are a typical sign that the administrator or the user lacks the required security awareness, putting at risk even a system that should be highly secure by its very design.

Viruses should not be confused with worms which belong to the world of networks entirely. Worms do not need a host to spread.

8.2.2 Network Security

Local security is concerned with keeping different users on one system apart from each other, especially from `'root'`. Network security, on the other hand, means that the system needs to be protected from an attack originating in the network.

The typical login procedure requiring a user name and a password for user authentication is a local security issue. However, in the particular case of logging in over a network, we need to differentiate between both security aspects. What happens until the actual authentication is network security and anything that happens afterwards is local security.

X Window System (X11 authentication)


As mentioned at the beginning, network transparency is one of the central characteristics of a UNIX system. X11, the windowing system of UNIX operating systems, can make use of this feature in an impressive way. With X11, it is basically no problem to log in at a remote host and start a graphical program that will then be sent over the network to be displayed on your computer. The protocol to communicate between the X application and the X server (which is the local process that draws the windows with the help of your video card) is relatively lightweight as far as bandwidth usage is concerned. This is because the protocol was designed in the eighties when network bandwidth was still a scarce resource.

Now if we want an X client to be displayed remotely using our X server, the latter is supposed to protect the resource managed by it (i. e. the display) from unauthorized access. In more concrete terms, certain permissions must be given to the client program. With the X Window System, there are two ways to do this, called host-based access control and cookie-based access control. The former relies on the IP address of the host where the client is supposed to run; the program to control this is `xhost`. What `xhost` does is to enter the IP address of a legitimate client into a tiny database belonging to the X server. Note, however, that relying on IP addresses for authentication is not very secure. For example, if there were a second user working on the host sending the client program, that user would have access to the X server as well — just like someone stealing the IP address. Because of these shortcomings, we will not describe this authentication method in more detail here, but you can learn about the way it functions if you read the man page of `xhost`, which includes a similar warning.

In the case of cookie-based access control, a character string is generated which is only known to the X server and to the legitimate user, just like an ID card of some kind. This cookie (the word goes back not to ordinary cookies, but to Chinese fortune cookies which contain an epigram) is stored on login in the file `.Xauthority` in the user's home directory and is available to any X Window client wanting to use the X server to display a window. The file `.Xauthority` can be examined by the user with the tool `xauth`. If you were to rename `.Xauthority` or if you deleted the file from your home directory by accident, you would not be able to open any new windows or X clients. Read

more about X Window security mechanisms in the man page of `xsecurity` (**man xsecurity**).

Apart from that, `ssh` (secure shell) can be used to completely encrypt a network connection and forward it to an X server transparently without the encryption mechanism being perceived by the user. This is also called X forwarding. X forwarding is achieved by simulating an X server on the server side and setting a `DISPLAY` variable for the shell on the remote host. Before being displayed, the client opens a connection with `sshd` (secure shell daemon, the server side program), which then gets the connections through to the real X server. If your setup requires that X clients are displayed remotely, consider using `ssh` and have a closer look at it. The man page of `ssh` has more information about the functionality of this program.



Caution

If you do not consider the host where you log in to be a secure host, do not use X forwarding. With X forwarding enabled, an attacker could authenticate via your `ssh` connection to intrude on your X server and sniff your keyboard input, for instance.

Buffer overflows and format string bugs

As discussed in the section on local security, buffer overflows and format string bugs should be classified as issues concerning both local and network security. As with the local variants of such bugs, buffer overflows in network programs, when successfully exploited, are mostly used to obtain `'root'` permissions. Even if that is not the case, an attacker could use the bug to gain access to an unprivileged local account to exploit any other vulnerabilities which might exist on the system.

Buffer overflows and format string bugs exploitable over a network link are certainly the most frequent form of remote attacks in general. Exploits for these — programs to exploit these newly-found security holes — are often posted on the security mailing lists. They can be used to target the vulnerability without knowing the details of the code. Over the years, experience has shown that the availability of exploit codes has contributed to more secure operating systems, obviously due to the fact that operating system makers were forced to fix the problems in their software. With free software, anyone has access to the source code (SuSE Linux comes with all available source codes) and anyone who finds a vulnerability and its exploit code can submit a patch to fix the corresponding bug.

DoS — Denial of Service

The purpose of this kind of attack is to force down a server program or even an entire system, something which could be achieved by various means: over-

loading the server, keeping it busy with garbage packets, or exploiting a remote buffer overflow.

Often a DoS attack is done with the sole purpose of making the service disappear. However, once a given service has become unavailable, communications could become vulnerable to so-called “man-in-the-middle attacks” (sniffing, TCP connection hijacking, spoofing) and DNS poisoning, explained below.

Man in the middle: sniffing, tcp connection hijacking, spoofing

In general, any remote attack performed by an attacker who puts himself between the communicating hosts is called a “man-in-the-middle attack”. What almost all types of man-in-the-middle attacks have in common is that the victim is usually not aware that there is something happening. There are many possible variants, for example, the attacker could pick up a connection request and forward that to the target machine himself. Now the victim has unwittingly established a connection with the wrong host, because the other end is posing as the legitimate destination machine. The simplest form of a man-in-the-middle attack is called “sniffer” — the attacker is “just” listening to the network traffic passing by. As a more complex attack, the “man in the middle” could try to take over an already established connection (hijacking). To do so, the attacker would have to analyze the packets for some time to be able to predict the TCP sequence numbers belonging to the connection. When the attacker finally seizes the role of the target host, the victims will notice this, because they get an error message saying the connection was terminated due to a failure.

What often makes things easier for attackers is the fact that there are protocols which are not secured against hijacking through encryption, but only perform a simple authentication procedure upon establishing the connection. Finally, we want to mention “spoofing”, an attack where packets are modified to contain counterfeit source data, mostly the IP address. Most active forms of attack rely on sending out such fake packets — something that, on a Linux machine, can only be done by the superuser (`'root'`).

Many of the attacks mentioned are carried out in combination with a DoS. If an attacker sees an opportunity to abruptly bring down a certain host, even if only for a short time, it will make it easier for him to push the active attack, because the host will not be able to interfere with the attack for some time.

DNS poisoning

DNS poisoning means that the attacker corrupts the cache of a DNS server by replying to it with spoofed DNS reply packets, trying to get the server to send certain data to a victim who is requesting information from that server. To foist such false information onto the server in a credible way, normally the attacker must have received and analyzed some packets from it. Given that many servers are configured to maintain a trust relationship with other hosts, based on IP addresses or host names, such an attack may be successful in a relatively short

time. On the other hand, it also requires quite an effort. In any case, the attacker will need a good understanding of the actual structure of the trust relationships between hosts. The attacker often needs to target a well-timed DoS attack at the name server, as well. Protect yourself by using encrypted connections that are able to verify the identity of the hosts to which to connect.

Worms

Worms are often confused with viruses, but there is a clear difference between the two. Unlike viruses, worms do not need to infect a host program to live. Rather, they are specialized to spread as quickly as possible on network structures. The worms that appeared in the past, such as Ramen, Lion, or Adore, make use of well-known security holes in server programs like `bind8` or `lprNG`. Protection against worms is relatively easy. Given that some time will elapse between the discovery of a security hole and the moment the worm hits your server, there is a good chance that an updated version of the affected program will be available on time. Of course, that is only useful if the administrator actually installs the security updates on the systems in question.

8.3 Some General Security Tips and Tricks

Information: To handle security competently, it is important to keep up with new developments and to stay informed about the latest security issues. One very good way to protect your systems against problems of all kinds is to get and install the updated packages recommended by security announcements as quickly as possible. SuSE security announcements are published on a mailing list to which you can subscribe by following the link <http://www.suse.de/security>. The list `suse-security-announce@suse.de` is a first-hand source of information regarding updated packages and includes members of SuSE's security team among its active contributors.

The mailing list `suse-security@suse.de` is a good place to discuss any security issues of interest. Subscribe to it under the URL as given above for `suse-security-announce@suse.de`.

`bugtraq@securityfocus.com` is one of the best-known security mailing lists worldwide. We recommend that you read this list, which receives between 15 and 20 postings per day. More information can be found at <http://www.securityfocus.com>.

The following is a list of rules which you may find useful in dealing with basic security concerns:

- According to the rule of using the most restrictive set of permissions possible for every job, avoid doing your regular jobs as `'root'`. This reduces the risk of getting a cuckoo egg or a virus and protects you from your own mistakes.
- If possible, always try to use encrypted connections to work on a remote machine. Use `"ssh"` (secure shell) to replace `telnet`, `ftp`, `rsh` and `rlogin`.

- Avoid using authentication methods based on IP addresses alone.
- Try to keep the most important network-related packages up-to-date and subscribe to the corresponding mailing lists to receive announcements on new versions of such programs (`bind`, `sendmail`, `ssh`, etc.). The same should apply to software relevant to local security.
- Change the `/etc/permissions` file to optimize the permissions of files crucial to your system's security. If you remove the `setuid` bit from a program, it might well be that it cannot do its job anymore in the way it is supposed to. On the other hand, consider that, in most cases, the program will also have ceased to be a potential security risk. You might take a similar approach with world-writable directories and files.
- Disable any network services you do not absolutely require for your server to work properly. This will make your system safer, plus it prevents your users from getting used to a service that you had never intended to be available in the first place (the so-called legacy problem). Open ports, with the socket state `LISTEN`, can be found with the program `netstat`. As for the options, we suggest that you use `netstat -ap` or `netstat -anp`. The `-p` option allows you to see which process is occupying a port under which name.

Compare the `netstat` results with those of a thorough port scan done from outside your host. An excellent program for this job is `nmap`, which not only checks out the ports of your machine, but also draws some conclusions as to which services are waiting behind them. However, port scanning may be interpreted as an aggressive act, so do not do this on a host without the explicit approval of the administrator. Finally, remember that it is important that you not only scan TCP ports, but also UDP ports (options `-ss` and `-su`).

- To monitor the integrity of the files of your system in a reliable way, use the program `tripwire`. Encrypt the database created by `tripwire` to prevent someone from tampering with it. Furthermore, keep a backup of this database available outside your machine, stored on an external data medium not connected to it by a network link.
- Take proper care when installing any third-party software. There have been cases where a hacker had built a trojan horse into the tar archive of a security software package, which was fortunately discovered very quickly. If you install a binary package, have no doubts about the site from which you downloaded it.

Note that SuSE's RPM packages are `gpg`-signed. The key used by SuSE for signing reads as follows:

```
ID:9C800ACA 2000-10-19 SuSE Package Signing Key <build@suse.de>
```

```
Key fingerprint = 79C1 79B2 E1C8 20C1 890F 9994 A84E DAE8 9C80 0ACA
```

The command `rpm -checksig package.rpm` shows whether the checksum and the signature of an uninstalled package are correct. Find the key on the first CD of the distribution and on most key servers worldwide.

- Check your backups of user and system files regularly. Consider that if you do not test whether the backup will work, it might actually be worthless.

- Check your log files. Whenever possible, write a small script to search for suspicious entries. Admittedly, this is not exactly a trivial task. In the end, only you can know which entries are unusual and which are not.
- Use `tcp_wrapper` to restrict access to the individual services running on your machine, so you have explicit control over which IP addresses can connect to a service. For further information regarding `tcp_wrappers`, consult the manual page of `tcpd(8)` and `hosts_access` (**man tcpd**, **man hosts_access**).
- Use SuSEfirewall to enhance the security provided by `tcpd` (`tcp_wrapper`). However, if you do not intend to provide any services from your host, you should probably install SuSE personal-firewall instead. Configuring SuSE personal-firewall is as simple as providing the name of the network interface on which inbound traffic should be rejected. Find more information on this in the files `/sbin/SuSEpersonal-firewall` and `/etc/rc.config.d/security.rc.config`.
- Design your security measures to be redundant: a message seen twice is much better than no message at all.

8.4 Using the Central Security Reporting Address

If you discover a security-related problem (please check the available update packages first), write an e-mail to security@suse.de. Please include a detailed description of the problem and the version number of the package concerned. SuSE will try to send a reply as soon as possible. You are encouraged to pgp encrypt your e-mail messages. SuSE's pgp key is as follows:

ID:3D25D3D9 1999-03-06 SuSE Security Team <security@suse.de>

Key fingerprint = 73 5F 2E 99 DF DB 94 C4 8F 5A A3 AE AF 22 F2 D5

This key is also available for download from: <http://www.suse.de/security>

9 Troubleshooting

9.1 Creating a Boot Disk

9.1.1 Creating a Boot Disk In DOS

Requirements

You need a formatted 3.5" floppy disk and a bootable 3.5" floppy drive. If you are working in Windows, launch setup from MS-DOS mode, *not* from inside a DOS window.

Additional information

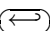
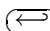
The disk images can be found on CD 1 in the directory `/disks`. These “images” can be copied to a floppy disk with the relevant utilities. The necessary disk images can be found on CD 1 in the directory `/disks`. These “images” can be copied to a floppy disk with the relevant utilities. CD 1, in the directory `/disks`, contains a number of disk images. Such an image can be copied to a disk with the help of suitable auxiliary programs. This disk is then called a boot disk. CD 1 in the directory `/disks` contains a number of disk images. Such an image can be copied to a disk with the help of suitable auxiliary programs, this disk is then called a boot disk. Included in these disk images are the “loader”, SYSLINUX, as well as the program `linuxrc`. SYSLINUX allows selection of a specific kernel for the booting process and to add parameters for your hardware, if necessary. The program `linuxrc` supports the loading of kernel modules for your hardware then starts the installation.

Normally the SuSE boot disk supplied can be used to boot. Only for exotic hardware not supported by the modularized kernel of this boot disk or if you download a disk image from the Internet (for example, from <ftp://ftp.suse.com>), do you need to create your own boot disk as described here.

With Setup

Step by step

To create a boot disk:

1. Start setup directly from CD 1.
2. Select ‘floppy’ and press . Next, select ‘Boot’ and press .

3. Select a disk with a suitable kernel, for example, that supports your SCSI adapter if you have one. `setup` shows the essential part of the kernel descriptions. If you need further information, look it up in `\disks\readme.dos`. Remember the name of your kernel. You will need it later. Now press `(←)`.
4. Create the boot disk. Insert the (DOS-formatted) disk into the 3.5" drive and select the disk to create.
 - Only the boot disk is needed ('Root' is not needed anymore for SuSE Linux.) Move the cursor onto 'Boot' and press `(←)`.
 - `setup` requests confirmation of disk insertion. Press `(←)` and the disk is written.
 - When this is finished, press `(←)`.
 - Now, select 'Done' to exit this screen and `setup`.

With rawrite

Alternatively, you might want to use the (perhaps slower) `rawrite`. If you are not running any other Unix/Linux system, you can use the DOS program `rawrite.exe` (CD 1, directory `\dosutils\rawrite`) to write the disk at the boot prompt.

The standard disk images are contained on CD 1 in the directory `/disks`. Read the file `README`. The image `bootdisk` or `scsi01` is the usual choice for the standard disk. Read the file there: `systypes.txt`. All the actual kernels can be found in the directory `/suse/images` (without extensions). Also read the `README` file there.

If you need the standard disk which is supplied with every SuSE Linux, e.g. the `aboot` floppy disk the standard disk for 32 bit systems proceed as follows. It is assumed that you are in the directory of the CD.

```
Q:> dosutils\rawrite\rawrite disks\bootdisk
```

You also need to create the "root disk", which contains the root file system needed by the installation tools.

If you need a specific type of support, another disk image should be used instead of `bootdisk`. If problems arise, `k_i386` can be implemented as a fallback kernel.

9.1.2 Creating a Boot Disk with UNIX

Requirements

You need access to a Unix or Linux system with an accessible CD-ROM drive and a formatted disk.

To create a boot disk:

1. If you need to format the disks first:


```
earth: # fdformat /dev/fd0u1440
```

2. Mount the first CD (disk 1) (e. g., to /cdrom):

```
earth: # mount -tiso9660 /dev/cdrom /cdrom
```

3. Change to the disks directory on CD:

```
earth: # cd /cdrom/disks
```

4. Create the boot disk with

```
earth: # dd if=/cdrom/disks/bootdisk of=/dev/fd0 bs=8k
```

```
earth: # dd if=/cdrom/disks/aboot_xx of=/dev/fd0 bs=8k
```

```
earth: # dd if=/cdrom/disks/bootdisk32 of=/dev/fd0 bs=8k
```

You also need to create the “root disk”, which contains the root file system needed by the installation tools.

Create this disk with the command:

```
earth: # dd if=/cdrom/disks/rootdisk32 of=/dev/fd0 bs=8k
```

In the README file in the directory disks,

In the file systypes.txt

read about what features specific kernels have. These files can be read with **more** or **less**.

5. If you need a different kernel, another disk image can be used in place of bootdisk. If problems arise, k_i386 can be implemented as a fallback kernel.

9.2 LILO Problems

Some Guidelines

Some simple guidelines at the beginning will avoid most LILO problems in advance (this is taken from the LILO user manual):

- *Don't panic!* If anything does not work, try to find the error or the cause first. Check the diagnosis before you start fixing the problem.
- Always have an up-to-date and tested *boot disk* at hand. SuSE Linux contains a full Linux system on its boot disk and installation CD (for the rescue system, see Section 9.3 on page 119) to allow you to reach all your Linux partitions. Tools are included for repairing almost any problems that can occur.
- Read the complete LILO documentation, especially if the system does not do what you want it to do.
- Check /etc/lilo.conf *before* using the map installer (/sbin/lilo).
- Be careful if you are using a large hard disk or multiple ones. Be aware of the 1024-cylinder limit.

- Try with and without the **linear** option (normally it should be better without).

9.2.1 Diagnosis of Errors: LILO Start Messages

This is mainly Section 5.2.1 from [Alm96].

When LILO loads itself, it displays the word 'LILO'. Each letter is printed before or after performing some specific action. If LILO fails at some point, the letters printed so far can be used to identify the problem.

nothing No part of LILO has been loaded. Either LILO is not installed at all or the partition on which its boot sector is located is not active.

'L' error ... The *first stage* boot loader has been loaded and started, but it cannot load the second stage boot loader (`/boot/boot.b`). The two-digit error codes indicate the type of problem. This condition usually indicates a media failure or a geometry mismatch.

'LI' The second stage has been invoked, but could not be started. This can either be caused by a geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LIL' The second stage of boot loader has been started, but it cannot load the descriptor table from the map file. This is typically due to a physical error of the boot device or a faulty disk geometry.

'LIL?' The second stage boot loader has been loaded at an incorrect address. This is typically caused by a subtle geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LIL-' The descriptor table in the map file is corrupt. This can either be caused by a geometry mismatch or by moving `/boot/boot.b` without reinstalling LILO.

'LILO' All parts of LILO have been successfully loaded.

Removing causes of error

The most common causes for *geometry errors* are not physical defects or invalid partition tables, but errors in LILO installation, including disregarding the 1024-cylinder limit (see next section) or an unsuccessful attempt at starting LILO from a logical partition.

In most cases, errors can be resolved using the following three methods:

1. Install the LILO data below the 1024-cylinder limit (if you have not already done so). This applies to the required Linux kernel, the directory contents of `/boot`, as well as the boot sector, which will incorporate the LILO start code.

2. Install LILO from scratch with **lilo** as 'root'. **lilo** will issue an informative log, if you increase its verbosity and create log files. This can be done with:

```
earth: # lilo -v -v -v >/boot/lilo.log 2>/boot/lilo.logerr
```

If the configuration is correct, `/boot/lilo.logerr` should be empty for a boot configuration. `/boot/lilo.log` includes exact information as to how LILO stores the locations of its files, which BIOS device numbers LILO uses for their respective hard disks, and more.

3. Check the consistency of the hard disk geometry data. In actuality, up to four areas of interest here:
 - a) The geometry LILO uses. See the log file mentioned above. Is influenced by the **disk** specification in `lilo.conf`
 - b) Geometry which was recognized by the Linux kernel. See the boot messages (`/var/log/boot.msg` or the command output **dmesg**). Informed by kernel parameters (to a certain extent).
 - c) Geometry which the partition table is based on. See the output of **fdisk -l**. Influenced by **fdisk** expert commands. Very risky for data! A full backup is, under any circumstances, highly recommended and really only for experts!
 - d) Geometry recognized by the BIOS. LILO discovers this geometry later when the system starts and is must be able to work with it. See the BIOS setup, and, if applicable, the SCSI host adapter (if available). This is influenced by the BIOS setup.

If there are discrepancies, making a good decision as to “where should I make adjustments?” is often the best method leading to the “path of least resistance”.

The following data should be examined when attempting to solve problems:

- `/etc/lilo.conf`
- Command output **fdisk -l** (partitioning)
- Above-mentioned log files
- BIOS and SCSI-BIOS hard disk settings

9.2.2 The 1024-Cylinder Limit




Note

Recently, BIOS versions are available which enable you to start operating systems above the 1024-cylinder limit. The current LILO version can use this BIOS extension. YaST and YaST2 will inform you accordingly of these options for your BIOS while configuring LILO. If your BIOS does not include this extension, continue reading here.

As emphasized before (on on page ??), the entire LILO machinery (including all data needed for booting) must be able to process BIOS calls, which means it must reside below the 1024--cylinder limit on the hard disk. The sections of the hard disk that can be used, called *allowed sections*, have already been discussed.

This restriction affects *only* the boot-up machinery. It is not required that LILO be installed on the Linux root partition. It is even possible, but quite dangerous, to put the boot machinery onto partitions of other operating systems to which Linux has read and write access.



Caution

Never install the LILO boot sector onto an unknown partition because you will severely damage the file system

The best method is to create a primary partition (within the allowed section) and to install all LILO files (including the LILO boot sector) into this partition. This will be, in most cases, the Linux root partition.

You can also add it to `/boot` with YaST. The only condition is that there has to be enough space for `boot.b`, `map`, `message`, and the Linux kernels that LILO should boot.

A few megabytes is enough. It does not matter where you put the rest of your partitions. There are no more restrictions. As soon as the kernel runs, you have unrestricted access to all installed drives.

But what to do if there is no space for such a partition? If you neither want to repartition your hard disk, upgrade to SCSI, or purchase a new BIOS version, there are still two (makeshift) possibilities:

- Use a boot disk instead of LILO on the hard disk or, if you are also running MS-DOS, use `loadlin`.
- Install the LILO boot machinery onto a Linux partition in the permitted section and where Linux has write access (e. g., a FAT or VFAT drive). We cannot put the LILO boot sector there as well. So there are only two places to put it. Either at the start of an extended partition on the first drive — as long as it is beneath the 1024--cylinder limit — or on the MBR.

Suppose that the partition in question is mounted on `/mnt`, that LILO is installed in the MBR (`/dev/hda`), and that you also boot DOS from `/dev/hda1`. Proceed as follows:

- Create a new directory (e. g., `/mnt/LINUX`) and copy the LILO files mentioned above to it (`boot.b`, `map`, `message`) as well as the chain loader of other operating systems (normally `chain.b`) and the Linux kernels that LILO should boot.
- Create a `/mnt/LINUX/lilo.conf` where all paths point to `/mnt/LINUX` (see File 9.2.1 on the next page).

```
# LILO Configuration file
# Start LILO global Section
boot=/dev/hda          # Installation target
backup=/mnt/LINUX/hda.xxxx # backup of old MBR
install=/mnt/LINUX/boot.b # Of course LILO and
map=/mnt/LINUX/map      # map file are in /mnt/LINUX!
message=/mnt/LINUX/message # optional
prompt
timeout=100           # Wait at prompt: 10 s
vga = normal          #
# End LILO global section
#
# Linux bootable partition config begins
image = /mnt/LINUX/First_Kernel # default
    root = /dev/Your_Root_Device # Root partition!
    label = linux
# Linux bootable partition config ends
#
# System section for other kernels:
#
# End Linux
# DOS bootable partition config begins
other = /dev/hda1      # MSDOS system drive
    label = dos
    loader = /mnt/LINUX/chain.b
    table = /dev/hda
# DOS bootable partition config ends
```

File 9.2.1: lilo.conf for Other Partitions

- Install LILO with *this* lilo.conf:

```
earth: # /sbin/lilo -C /mnt/LINUX/lilo.conf
```

After that, LILO should work. Boot MS-DOS and protect the LILO files as well as possible against write access (any write access disables LILO). To accomplish this, assign to all files in X:\LINUX (where the 'X' is the DOS drive mounted to /mnt) the DOS attributes *system* and *hide*.

In conclusion, we point you toward two HOWTOs in /usr/share/doc/howto/en/mini/ — LILO.gz and Large-Disk.gz.

9.3 The SuSE Rescue System

The rescue system is launched using the SuSE boot disk or from your bootable SuSE Linux CD 1. It is required that the disk and CD-ROM drives are bootable. If necessary, you will need to change the boot series in the CMOS setup.

Following are the steps for starting the rescue system:

1. Start your system with the SuSE boot disk or with the first SuSE Linux CD inserted in your CD-ROM drive.

2. Launch the entire system or, at the boot prompt, either enter `yast` or `manual`, where you can define which kernel modules should be loaded.
3. Make the respective settings for language, keyboard, and screen.
4. Select the item 'Installation/Start system' in the main menu.
5. If you started with the boot disk, you should now insert the installation CD or the rescue disk with the compressed image of the rescue system.
6. In the menu 'Start installation/system' select the item 'Start rescue system' then specify the desired source medium. Subsequently, we will introduce a few tips on selection options:

'CD-ROM': When loading the rescue system, the path `/cdrom` is exported. This makes the installation from *this* CD possible.

Note

You now still need to enter the required values in `SuSEconfig`.



'Network (NFS)': To start the `rescue` system via NFS from the network, you must have the driver for your network card already installed.

'Network (FTP)': To start the `rescue` system via FTP from the network, you have to have your network card driver ready.

'hard disk': Load the `rescue` system from the hard disk.

'Floppy Disk': The `rescue` system can also be started from the floppy disk, especially if the computer only has a small amount of working memory.

Regardless of the medium chosen, the rescue system will be decompressed, loaded onto a RAM floppy disk as a new root file system, mounted, and started. Now it is ready for use.

9.3.1 Working with the Rescue System

The rescue system provides three virtual consoles on keys `(Alt) + (F1)` to `(Alt) + (F3)`. Here `'root'` may log in without a password. `(Alt) + (F4)` accesses the system console where you can view the kernel and syslog messages.

A shell and lots of other useful utilities (net tools), such as the `mount` program, can be found in the `/bin` directory. In `sbin`, find important file and network utilities for reviewing and repairing the file system (e.g., `e2fsck`).

Furthermore, this directory contains the most important binaries for system maintenance, such as `fdisk`, `mkfs`, `mkswap`, `mount`, `mount`, `init`, and `shutdown`, as well as `ifconfig`, `route`, and `netstat` for maintaining the network.

An editor, `vi`, is located in `/usr/bin`. Also, tools like `grep`, `find`, and `less`, along with `telnet` are available.

Accessing Your Normal System

To mount your SuSE Linux system using the rescue system, use the mountpoint `/mnt`. You can also use or create another directory.


As an example, assume your normal system is put together according to the `/etc/fstab` shown in the example File 9.3.1.

<code>/dev/sdb5</code>	<code>swap</code>	<code>swap</code>	<code>defaults</code>	<code>0</code>	<code>0</code>
<code>/dev/sdb3</code>	<code>/</code>	<code>ext2</code>	<code>defaults</code>	<code>1</code>	<code>1</code>
<code>/dev/sdb6</code>	<code>/usr</code>	<code>ext2</code>	<code>defaults</code>	<code>1</code>	<code>2</code>

File 9.3.1: Example `/etc/fstab`

Caution

Pay attention to the order of steps outlined in the following section for mounting the various devices.



To access your entire system, mount it step-by-step in the `/mnt` directory using the following commands:

```
earth:/ # mount /dev/sdb3 /mnt
earth:/ # mount /dev/sdb6 /mnt/usr
```

Now you can access your entire system and, for example, correct mistakes in configuration files such as `/etc/fstab`, `/etc/passwd`, and `/etc/inittab`. The configuration files are now located in the `/mnt/etc` directory instead of in `/etc`.

To recover even completely lost partitions with the `fdisk` program by simply setting it up again, determine where on the hard disk the partitions were previously located and make a hardcopy printout of the `/etc/fstab` directory as well as the output of the command

```
earth: # fdisk -l /dev/<disk>
```

Instead of the `<disk>` variable, insert, in order, the device names of your hard disks, i.e., `hda`.

Repairing File Systems

Damaged file systems are tricky problems for the rescue system. This could happen after an unscheduled shutdown caused by power failure or a system crash. Generally, file systems cannot be repaired on a running system. If you encounter really severe problems, you may not even be able to mount your root file system and have the system boot end in a "kernel panic". Here, the only chance is to repair the system from the "outside" using a rescue system.

The SuSE Linux rescue system contains the utilities `e2fsck` and `dumpe2fs` (for diagnosis). These should remedy most problems. In an emergency, man pages often are not available. That is why we have included them in this manual in Appendix 9.3.1 on page 123.

Example: If mounting a file system fails due to an *invalid* superblock, the `e2fsck` program would probably fail, too. If this were the case, your superblock may be corrupted, too. There are copies of the superblock located every 8192 blocks (8193, 16385, etc.). If your superblock is corrupted, try one of the copies instead. This is accomplished by entering the command:

```
earth: # e2fsck -f -b 8193 /dev/damaged_partition
```

The `-f` option forces the file system check and overrides `e2fsck`'s error so that — since the superblock copy is intact — everything is fine.

e2fsck Manual Page

E2FSCK(8)

E2FSCK(8)

NAME

e2fsck - check a Linux second extended file system

SYNOPSIS

```
e2fsck [ -pacnyrdfvstFSV ] [ -b superblock ] [ -B block-size ] [ -l|-L bad_blocks_file ] [ -C fd ] [ -j external-journal ] [ device ]
```

DESCRIPTION

e2fsck is used to check a Linux second extended file system (e2fs). E2fsck also supports ext2 file systems containing a journal, which are also sometimes known as ext3 file systems.

device is the special file corresponding to the device (e.g /dev/hdcl).

OPTIONS

-a This option does the same thing as the -p option. It is provided for backwards compatibility only; it is suggested that people use -p option whenever possible.

-b superblock

Instead of using the normal superblock, use an alternative superblock specified by superblock. This option is normally used when the primary superblock has been corrupted. The location of the backup superblock is dependent on the file system's block size. For file systems with 1k block sizes, a backup superblock can be found at block 8193; for file systems with 2k block sizes, at block 16384; and for 4k block sizes, at block 32768.

Additional backup superblock can be determined by using the mke2fs program using the -n option to print out where the superblock were created. The -b option to mke2fs, which specifies block size of the file system must be specified in order for the superblock locations that are printed out to be accurate.

If an alternative superblock is specified and the file system is not opened read-only, e2fsck will make sure that the primary superblock is updated appropriately upon completion of the file system check.

-B blocksize

Normally, e2fsck will search for the superblock at various different block sizes in an attempt to find the appropriate block size. This search can be fooled in some cases. This option forces e2fsck to only try locating the superblock at a particular block size. If the superblock is not found, e2fsck will terminate with a fatal error.

- c This option causes e2fsck to run the badblocks(8) program to find any blocks which are bad on the file system, and then marks them as bad by adding them to the bad block inode.
- C This option causes e2fsck to write completion information to the specified file descriptor so that the progress of the file system check can be monitored. This option is typically used by programs which are running e2fsck. If the file descriptor specified is 0, e2fsck will print a completion bar as it goes about its business. This requires that e2fsck is running on a video console or terminal.
- d Print debugging output (useless unless you are debugging e2fsck).
- f Force checking even if the file system seems clean.
- F Flush the file system device's buffer caches before beginning. Only really useful for doing e2fsck time trials.
- j external-journal
Set the pathname where the external-journal for this file system can be found.
- l filename
Add the blocks listed in the file specified by filename to the list of bad blocks. The format of this file is the same as the one generated by the badblocks(8) program.
- L filename
Set the bad blocks list to be the list of blocks specified by filename. (This option is the same as the -l option, except the bad blocks list is cleared before the blocks listed in the file are added to the bad blocks list.)
- n Open the file system read-only, and assume an answer of 'No' to all questions. Allows e2fsck to be used non-interactively. (Note: if the -c, -l, or -L options are specified in addition to the -n option, then the file system will be opened read-write, to permit the bad-blocks list to be updated. However, no other changes will be made to the file system.)
- p Automatically repair ("preen") the file system without any questions.
- r This option does nothing at all; it is provided only for backwards compatibility.
- s This option will byte-swap the file system so that it is using the normalized, standard byte-order (which is i386 or little endian). If the filesystem is already in the standard byte-order, e2fsck will take no action.

- S This option will byte-swap the filesystem, regardless of its current byte-order.
- t Print timing statistics for e2fsck. If this option is used twice, additional timing statistics are printed on a pass by pass basis.
- v Verbose mode.
- V Print version information and exit.
- y Assume an answer of 'Yes' to all questions; allows e2fsck to be used non-interactively.

EXIT CODE

The exit code returned by e2fsck is the sum of the following conditions:

- 0 - No errors
- 1 - File system errors corrected
- 2 - File system errors corrected, system should be rebooted if file system was mounted
- 4 - File system errors left uncorrected
- 8 - Operational error
- 16 - Usage or syntax error
- 128 - Shared library error

SIGNALS

The following signals have the following effect when sent to e2fsck.

SIGUSR1

This signal causes e2fsck to start displaying a completion bar. (See discussion of the -C option.)

SIGUSR2

This signal causes e2fsck to stop displaying a completion bar.

REPORTING BUGS

Almost any piece of software will have bugs. If you manage to find a filesystem which causes e2fsck to crash, or which e2fsck is unable to repair, please report it to the author.

Please include as much information as possible in your bug report. Ideally, include a complete transcript of the e2fsck run, so I can see exactly what error messages are displayed. If you have a writeable filesystem where the transcript can be stored, the script(1) program is a handy way to save the output of e2fsck to a file.

It is also useful to send the output of dumpe2fs(8). If a specific inode or inodes seems to be giving e2fsck trouble, try running the debugfs(8) command and send the output of the stat(1u) command run on the relevant inode(s). If the inode is a directory, the debugfs dump command will allow you to extract the contents of the directory inode, which can sent to me after being first run through uenocode(1).

Always include the full version string which e2fsck displays when it is run, so I know which version you are running.

AUTHOR

This version of e2fsck was written by Theodore Ts'o
<tytso@mit.edu>.

SEE ALSO

mke2fs(8), tune2fs(8), dumpe2fs(8), debugfs(8)

E2fsprogs version 1.23 August 2001

3

Bibliography

- [Alm96] ALMESBERGER, Werner: *LILO User's guide*, 1996. – (see file `/usr/share/doc/lilo/user.dvi`)
- [Bai97] BAILEY, Edward C.: *Maximum RPM*. Red Hat, 1997. – (ISBN 1-888172-78-9)
- [BBD⁺97] BECK, Michael; BÖHME, Harald; DZIADZKA, Mirko; KUNITZ, Ulrich; MAGNUS, Robert ; VERWORNER, Dirk: *Linux-Kernel-Programmierung*. 4th edition. Addison Wesley GmbH, 1997. – (ISBN 3-8273-1144-6)
- [BD98] BORKNER-DELCARLO, Olaf: *Linux im kommerziellen Einsatz*. Carl Hanser Verlag, 1998. – (ISBN 3-446-19465-7)
- [BD00] BORKNER-DELCARLO, Olaf: *Das Samba Buch*. 2nd edition. SuSE PRESS, 2000. – (ISBN 3-934678-22-X)
- [CAR93] COSTALES, Bryan; ALLMAN, Eric ; RICKERT, Neil: *sendmail*. O'Reilly & Associates, Inc., 1993. – (ISBN 1-56592-056-2)
- [CB96] CHESWICK, William R.; BELLOVIN, Steven M.: *Firewalls und Sicherheit im Internet*. Addison Wesley GmbH, 1996. – (ISBN 3-89319-875-x)
- [CR91] CAMERON, Debra; ROSENBLATT, Bill: *Learning GNU Emacs*. O'Reilly & Associates, Inc., 1991. – (ISBN 0 937175-84-6)
- [CZ96] CHAPMAN, Brent; ZWICKY, Elisabeth D.: *Einrichten von Internet Firewalls. Sicherheit im Internet gewährleisten..* O'Reilly & Associates, Inc., 1996. – (ISBN 3-930673312)
- [Deu01] DEUTSCH, Karl: *SuSE Linux: System und Anwendungen im Überblick*. SuSE PRESS, 2001. – (ISBN 3-934678-41-6)
- [DR98] DAWSON, Terry; RUBINI, Alessandro: *NET-3 HOWTO*, v1.4, August 1998. – (see file `/usr/share/doc/howto/en/NET-3-HOWTO.gz`)
- [DR99] DAWSON, Terry; RUBINI, Alessandro: *NET3-4 HOWTO*, v1.5, August 1999. – (see file `/usr/share/doc/howto/en/NET3-4-HOWTO.gz`)
- [EH98] ECKEL, George; HARE, Chris: *Linux – Internet Server*. Carl Hanser Verlag, 1998. – (ISBN 3-446-19044-9)
- [Fri93] FRISCH, Aileen: *Essential System Administration*. O'Reilly & Associates, Inc., 1993. – (ISBN 0-937175-80-3)
- [FW00] FISHER, Stefan; WALTHER, Ulrich: *Linux Netzwerke*. SuSE PRESS, 2000. – (ISBN 3-934678-20-3)
- [Gil92] GILLY, Daniel: *UNIX in a nutshell: System V Edition*. O'Reilly & Associates, Inc., 1992. – (ISBN 1-56592-001-5)

- [Gri94] GRIEGER, W.: *Wer hat Angst vorm Emacs?*. Addison Wesley GmbH, 1994. – (ISBN 3-89319-620-X)
- [GS93] GARFINKEL, Simson; SPAFFORD, Gene: *Practical UNIX Security*. O'Reilly & Associates, Inc., 1993. – (ISBN 0-937175-72-2)
- [Hei96] HEIN, Jochen: *Linux-Companion zur Systemadministration*. Addison Wesley GmbH, 1996. – (ISBN 3-89319-869-5)
- [Her92] HEROLD, H.: *UNIX Grundlagen*. Addison Wesley GmbH, 1992. – (ISBN 3-89319-542-8)
- [HHMK96] HETZE, Sebastian; HOHNDEL, Dirk; MÜLLER, Martin ; KIRCH, Olaf: *Linux Anwenderhandbuch*. 6th edition. LunetIX Softfair, 1996. – (ISBN 3-929764-05-9)
- [Hof97] HOFFMANN, Erwin: EMail-Gateway mit qmail. In: *iX* 12 (1997), S. 108ff.
- [HR98] HÖLZER, Matthias; RÖHRIG, Bernhard: *KDE – Das K Desktop Environment*. Computer & Literatur, 1998. – (ISBN 3-932311-50-7)
- [HST97] HOLZ, Helmut; SCHMITT, Bernd ; TIKART, Andreas: *Linux für Internet & Intranet*. International Thomson Publishing, 1997. – (ISBN 3-8266-0342-7)
- [Hun95] HUNT, Craig: *TCP/IP Netzwerk Administration*. O'Reilly & Associates, Inc., 1995. – (ISBN 3-930673-02-9)
- [JT98] JOHNSON, Michael K.; TROAN, Erik W.: *Anwendungen entwickeln unter Linux*. Addison Wesley GmbH, 1998. – (ISBN 3-8273-1449-6)
- [Kie95] KIENLE, Micheal: TIS: Toolkit für anwendungsorientierte Firewall-Systeme. In: *iX* 8 (1995), S. 140ff.
- [Kir95] KIRCH, Olaf: *LINUX Network Administrator's Guide*. O'Reilly & Associates, Inc., 1995. – (ISBN 1-56592-087-2)
- [Kof99] KOFLER, Michael: *Linux – Installation, Konfiguration, Anwendung*. 4th edition. Addison Wesley GmbH, 1999. – (ISBN 3-8273-1475-5)
- [Kri00] KRIENKE, Reiner: *Kommunikation unter Linux*. SuSE PRESS, 2000. – (ISBN 3-934678-23-8)
- [Kun95] KUNITZ, Ulrich: Sicherheit fast kostenlos: Einrichtung eines kostenlosen Firewall-Systems. In: *iX* 9 (1995), S. 176ff.
- [Lam90] LAMB, Linda: *Learning the vi Editor*. O'Reilly & Associates, Inc., 1990. – (ISBN 0-937175-67-6)
- [Lef96] LEFFLER, Sam: *HylaFAX Home Page*, 1996
- [Moh98] MOHR, James: *UNIX-Windows-Integration*. International Thomson Publishing, 1998. – (ISBN 3-8266-4032-2)
- [OT92] O'REILLY, Tim; TODINO, Grace: *Managing UUCP and Usenet*. O'Reilly & Associates, Inc., 1992. – (ISBN 0-937175-93-5)
- [Per94] PERLMAN, G.: *Unix For Software Developers*. Prentice-Hall, 1994. – (ISBN 13-932997-8)

- [POL97] PEEK, Jerry; O'REILLY, Tim ; LOUKIDES, Mike: *Unix Power Tools*. 2nd edition. Sebastopol : O'Reilly & Associates, Inc., 1997
- [Pug94] PUGH, K.: *UNIX For The MS-DOS User*. Prentice-Hall, 1994. – (ISBN 13-146077-3)
- [Rub98] RUBINI, Alessandro: *Linux-Gerätetreiber*. O'Reilly & Associates, Inc., 1998. – (ISBN 3-89721-122-X)
- [SB92] SCHOONOVER, M.; BOWIE, J.: *GNU Emacs*. Addison Wesley GmbH, 1992. – (ISBN 0-201-56345-2)
- [Sch98] SCHEIDERER, Jürgen: Sicherheit Kostenlos - Firewall mit Linux. In: *iX* 12 (1998)
- [Sto98] STOLL, Clifford: *Kuckucksei; Die Jagd auf die deutschen Hacker, die das Pentagon knackten*. Fischer-TB.-Vlg., 1998. – (ISBN 3596139848)
- [The96] THE XFREE86™-TEAM: *XF86Config(4/5) - Configuration File for Xfree86™*, 1996. – Manual-Page zu XFree86™
- [TSP93] TODINO, Grace; STRANG, John ; PEEK, Jerry: *Learning the UNIX operating system*. O'Reilly & Associates, Inc., 1993. – (ISBN 1-56592-060-0)
- [Wel94] WELSH, Matt: *Linux Installation and Getting Started*. 2. Aufl. SuSE GmbH, 1994. – (ISBN 3-930419-03-3)
- [WK95] WELSH, Matt; KAUFMAN, Lars: *Running Linux*. O'Reilly & Associates, Inc., 1995. – (ISBN 1-56592-100-3)
- [WK98] WELSH, Matt; KAUFMAN, Lars: *Linux – Wegweiser zur Installation & Konfiguration*. 2nd edition. O'Reilly & Associates, Inc., 1998. – (ISBN 3-930673-58-4)

Index

symbols

/etc/init.d/nfsserver 86
 /etc/init.d/portmap 86
 \$ALICE_HOME 98
 1024 Zylinder
 LILO problems 117

A

access to file server
 with MacOS 75
 with Windows 2000 72
 with Windows 95/98/ME .. 69
 administrator account 26
 alice 97
 ALICE 100
 Andrew Tridgell 87
 Apache 91–93
 applications 3

B

backups
 creating with YaST 63
 bash 86
 bind 81
 BIND 82
 bind8 83
 boot disk
 creating with Unix 114
 boot disks 61
 creating in DOS 113
 creating with rawrite 114
 creating with setup 113
 boot manager 24
 boot modes 60
 boot settings 54
 booting 123

C

Cache Manager 91
 Calamaris 92
 check 123
 client
 configuration 69

configuration files

lilo.conf 104
 permissions 111
 permissions.local 105
 rc.config 105
 configuring
 YaST2 35
 crash 123

D

DHCP 83
 DNS 81
 security and 109
 Domain Name Service *see* DNS
 drivers
 installing with YaST2 63
 drives
 link 72
 dumpe2fs 121
 Dynamic Host Configuration
 Protocol *see* DHCP

E

e-mail
 configuring on network with
 YaST2 51
 e2fsck 120–122
 Manual Page 123
 emergency system 119
 exports 86
 exports 48

F

fdisk 120, 121
 file service 85
 file systems
 repairing 121
 find 120
 firewall 44, 88
 firewalls
 YaST2 and 41
 free software 5

G

GDI printers 38
 GNU General Public License ...
 see GPL
 GPL 5
 grep 120
 groups
 managing with YaST2 53

H

hardware
 information 62
 YaST2 and 36
 host 83
 host name 32

I

ifconfig 120
 init 120
 installation sources 57
 Internet
 connection configuration in
 YaST2 40
 networks and 42
 Internet access 76
 access controls 77
 Intranet server 92
 IP addresses 40
 ISDN
 YaST2 and 43

K

kernel 1
 keyboard
 configuring with YaST2 ... 40
 kinternet 43

L

less 120
 LILO 24–26, 28, 60, 61,
 115–119
 problems
 1024-cylinder limit 117

132

- configuring in YaST2 39
- security 107
- X11 39
- XFree86 ii
- Y**
- YaST ii, 58, 63, 64, 67, 99, 100, 117, 118
 - backups 63
 - driver installation 63
- YaST1 105
- YaST2 .. 2, 7, 11–17, 19–25, 27, 30, 31, 35–39, 42, 44–46, 48–67, 82, 83, 89, 105, 117
 - boot modes 60
 - crypto file system 23
 - encryption 23
 - firewalls 41
 - group management 53
 - hardware 36
 - hardware information 62
 - installation source 57
 - Internet configuration 40
 - ISDN 43
 - keyboard 40
 - Lilo 24
 - Logical Volume Manager . 19
 - LVM 19
 - modem 45
 - network card 45
 - NFS server 48
 - NIS server 49
 - Online Update 58
 - Online Update in the console . 58
 - printer 37
 - rc.config editor 64
 - Samba 47
 - security 54
 - sendmail 51
 - SLCS variables 64
 - software installation 56
 - software removal 56
 - update 59
 - user management 52
 - X configuration 39
- YOU 58